

# Space Lower Bounds for Itemset Frequency Sketches\*

Edo Liberty<sup>†</sup>   Michael Mitzenmacher<sup>‡</sup>   Justin Thaler<sup>†</sup>   Jonathan Ullman<sup>§</sup>

## Abstract

Given a database, computing the fraction of rows that contain a query itemset or determining whether this fraction is above some threshold are fundamental operations in data mining. A uniform sample of rows is a good sketch of the database in the sense that all sufficiently frequent itemsets and their approximate frequencies are recoverable from the sample, and the sketch size is independent of the number of rows in the original database. For many seemingly similar problems there are better sketching algorithms than uniform sampling. In this paper we show that for itemset frequency sketching this is not the case. That is, we prove that there exist classes of databases for which uniform sampling is a space optimal sketch for approximate itemset frequency analysis, up to constant or iterated-logarithmic factors.

## 1 Introduction

Identifying frequent itemsets is one of the most basic and well-studied problems in data mining. Formally, we are given a binary database  $\mathcal{D} \in (\{0, 1\}^d)^n$  consisting of  $n$  rows and  $d$  columns, or attributes.<sup>1</sup> An *itemset*  $T \subseteq [d]$  is a subset of the attributes, and the *frequency*  $f_T$  of  $T$  is the fraction of rows of  $\mathcal{D}$  that have a 1 in all columns in  $T$ .

Computing itemset frequencies is a central primitive that can be used to solve the following problems (and countless others): given a large corpus of text files, compute the number of documents containing a specific search query; given user records, compute the fraction of users who belong to a specific demographic; given event logs, compute sets of events that are observed together; given shopping cart data, identify bundles of items that are frequently bought together.

In many settings, an approximation of  $f_T$ , as opposed to an exact result, suffices. Alternatively, in some settings it suffices to recover a single bit indicating whether or not  $f_T \geq \epsilon$  for some user defined threshold  $\epsilon$ ; such frequent itemsets may require additional study or processing. It is easy to show that uniformly sampling  $\text{poly}(d/\epsilon)$  rows from  $\mathcal{D}$  and computing the approximate frequencies on the sample  $\mathcal{S}(\mathcal{D})$  provides good approximations to  $f_T$  up to additive error  $\epsilon$ . Our main contribution is to provide lower bounds establishing that uniform sampling is an essentially optimal *sketch*, in terms of the space/accuracy tradeoff, for many parameter regimes. Here, a sketch  $\mathcal{S}(\mathcal{D})$  of a database is a bit string that enables recovery of accurate approximations to itemset frequencies.

Note that, unlike a row sample, in general a sketch is not limited to containing a subset of the database rows. Our lower bounds hold for any summary data structure and recovery algorithm that constitute a valid sketch.

---

\*This paper supersedes an earlier manuscript of the same title by the first three authors.

<sup>†</sup>Yahoo Labs

<sup>‡</sup>Harvard University, School of Engineering and Applied Sciences.

<sup>§</sup>Northeastern University, College of Computer and Information Sciences.

<sup>1</sup>Throughout, we use the terms *attributes* and *items* interchangeably. While attributes may be non-binary in many applications, any attribute with  $m$  possible values can be decomposed into  $2\lceil \log m \rceil$  binary attributes, using two binary attributes to mark whether the value is 0 or 1 in the  $i$ th bit. We therefore focus on the binary case.

## 1.1 Motivation

### 1.1.1 The Case Against Computing Frequent Itemsets Exactly

If the task is only to identify frequent itemsets ( $f_T \geq \epsilon$  for some  $\epsilon$ ), it is natural to ask whether we can compute all  $\epsilon$ -frequent itemsets and store only those. Assuming that only a small fraction of itemsets are  $\epsilon$ -frequent, this will result in significant space saving relative to naïve solutions. The extensive literature on exact mining of frequent itemsets dates back to work of Agrawal et al. [AIS93], whose motivation stemmed from the field of market basket analysis. As the search space for frequent itemsets is exponentially large (i.e., size  $2^d$ ), substantial effort was devoted to developing algorithms that rapidly prune the search space and minimize the number of scans through the database. While the algorithms developed in this literature offer substantial concrete speedups over naïve approaches to frequent itemset mining, they may still take time  $2^{\Omega(d)}$ , simply because there may be these many frequent itemsets. For example, if there is a frequent itemset of cardinality  $d/10$ , each of its  $2^{d/10}$  subsets is also frequent. Motivated by this simple observation, there is now an extensive literature on condensed or non-redundant representations of exact frequent itemsets. Reporting only *maximal* and *closed* frequent itemsets often works well in practice, but it still requires exponential size in the worst case (see the survey [CG07]).

Irrespective of space complexity, the above methods face computational challenges. Yang [Yan04] determined that counting all frequent itemsets is #P-complete, and a bottleneck for enumeration is that the number of frequent itemsets can be exponentially large. Hamilton et al. [HCW06] provide further hardness results based on parametrized complexity. Here we observe that finding even a single frequent itemset of approximately maximal size is NP-hard. (The authors of [LLSW05] noticed this connection as well but did not mention approximation-hardness.)

Consider the bipartite graph containing  $n$  nodes (rows) on one side and  $d$  nodes (attributes) on the other. An edge exists between the two sides if and only if the row contains the attribute with value 1. Assume there exists a frequent itemset of cardinality  $\epsilon n$  and frequency  $\epsilon$ . This itemset induces a balanced complete bipartite subgraph with  $\epsilon n$  nodes on each side. Likewise, any balanced complete bipartite subgraph with  $\epsilon n$  nodes per side implies the existence of an itemset of cardinality  $\epsilon n$  and frequency  $\epsilon$ . Finding the maximal balanced complete bipartite subgraph is NP-hard, and approximating it requires superpolynomial time assuming that SAT cannot be solved in subexponential time [FK04]. Hence, finding an itemset of approximately maximal frequency requires superpolynomial time under the same assumption.

### 1.1.2 The Case for Itemset Sketches

Determining the smallest possible itemset sketches (as defined in § 1.3) is of interest in several data analysis settings.

**Interactive Knowledge Discovery.** Knowledge discovery in databases is often an interactive process: an analyst poses a sequence of queries to the dataset, with later queries depending on the answers to earlier ones [MT96]. For large databases, it may be inefficient or even infeasible to reread the entire dataset every time a query is posed. Instead, a user can keep around an itemset sketch only; this sketch will be much smaller than the original database, while still providing fast and accurate answers to itemset frequency queries.

**Efficient Data Release.** Itemset oracles capture a central problem in *data release*. Here, a data curator (such as a government agency like the US Census Bureau) wants to make a dataset publicly available. Due to their utility and ease of interpretation, the data format of choice in these settings is typically *marginal contingency tables* (marginal tables for short). For any itemset  $T \subseteq [d]$  with  $|T| = k$ , the marginal table corresponding to  $T$  has  $2^k$  entries, one for each possible setting of the attributes in  $T$ ; each entry counts

how many rows in the database are consistent with the corresponding setting of the  $k$  attributes. Notice that marginal tables are essentially just a list of itemset frequencies for  $\mathcal{D}$ .<sup>2</sup>

However, marginal tables can be extremely large (as any  $k$ -attribute marginal table has  $2^k$  entries and there  $\binom{d}{k}$  such tables), and each released table may be downloaded by an enormous number of users. Rather than releasing marginal tables in their entirety, the data curator can instead choose to release an itemset summary. This summary can be much smaller than even a single  $k$ -attribute marginal table, while still permitting any user to obtain fast and accurate estimates for the frequency of any  $k$ -attribute marginal query.

**Mitigating Runtime Bottlenecks.** While the use of itemset sketches cannot circumvent the hardness results discussed in Section 1.1.1, in many settings the empirical runtime bottleneck is the number of scans through the database, rather than the total runtime of the algorithm. The use of itemset sketches eliminates the need for the user to repeatedly scan or even keep a copy of the original database. The user can instead run a computationally intensive algorithm on the sketch to solve (natural approximation variants) of the hard decision or search problems. Indeed, there has been considerable work in the data mining community devoted to bounding the magnitude of errors that build up as a result of using approximate itemset frequency information when performing more complicated data mining tasks, such as rule identification [MT96].

## 1.2 Other Prior Work

The idea of producing condensed representations of approximate frequent itemsets is not new. Most relevant to our work, an influential paper by Mannila and Toivonen defined the notion of an  $\epsilon$ -adequate representation of any class of queries [MT96]. Our For-All-Itemset-Frequency-Estimator sketching task essentially asks for an  $\epsilon$ -adequate representation for the class of all itemset frequency queries. Mannila and Toivonen analyzed the magnitude of errors that build up when using  $\epsilon$ -adequate representations to perform more complicated data mining tasks, such as rule identification. Subsequent work by Boulicaut et al. [BBR03] presented algorithms yielding  $\epsilon$ -adequate representations for the class of all itemset queries, while Pei et al. [PDZH04] gave algorithms for approximating the frequency of all *frequent* itemsets to error  $\epsilon$ . Unlike the trivial algorithms that we describe in Section 2, the algorithms presented in [MT96, BBR03, PDZH04] take exponential time in the worst case, and do not come with worst-case guarantees on the size of the output.

Streaming algorithms for both exact and approximate variants of frequent itemset mining have also been extensively studied, in a line of work dating back to Manku and Motwani [MM02] — see the survey [CKN08]. None of these works have been able to show that these algorithms use less space than uniform random sampling of database rows, and our results justify why. In particular, to our knowledge there has been no work establishing lower bounds on the space complexity of streaming algorithms for identifying approximate frequent itemsets that are better than the lower bounds that hold for the much simpler *approximate frequent items* problem (a.k.a. the heavy hitters problem). Note that the lower bounds that we establish in this work apply even to summaries computed by non-streaming algorithms.

**Related Work by Price [Pri14].** An earlier version of this manuscript by the first three authors [LMT14] gave lower bounds on the size of sketches for frequent itemset mining that are quantitatively weaker than the bounds presented here [LMT14]. In work subsequent to [LMT14], and contemporaneous with the work presented here, Price [Pri14] discovered a short proof of an optimal lower bound for the For-All-Itemset-Frequency-Indicator sketching problem (defined in Section 1.3 below) for itemsets  $T$  of size  $|T| = O(1)$ . We compare our results to Price’s in more detail in Section 3.1.

---

<sup>2</sup>More precisely, itemset frequency queries are equivalent to *monotone conjunction* queries on a database, while marginal tables are equivalent to general (non-monotone) conjunction queries.

### 1.3 Notation and Problem Statements

Throughout,  $\mathcal{D} \in (\{0, 1\}^d)^n$  will denote a binary database consisting of  $n$  rows and  $d$  columns, or attributes. We denote the set  $\{1, \dots, d\}$  by  $[d]$ . An *itemset*  $T \subseteq [d]$  is a subset of the attributes; abusing notation, we also use  $T$  to refer to the *indicator vector* in  $\{0, 1\}^d$  whose  $i$ th entry is 1 if and only if  $i \in T$ . We refer to any itemset  $T$  with  $|T| = k$  as a *k-itemset*. The  $i$ th row of  $\mathcal{D}$  will be denoted by  $\mathcal{D}(i)$ , and the  $j$ th entry of the  $i$ th row will be denoted by  $\mathcal{D}(i, j)$ . We say that a row *contains* an itemset  $T$  if the row has a 1 in all columns in  $T$ . The *frequency*  $f_T(\mathcal{D})$  of  $T$  is the fraction of rows of  $\mathcal{D}$  that contain  $T$ . Alternatively,  $f_T(\mathcal{D}) = \frac{1}{n} \sum_{i=1}^n \mathbb{I}_{\{T \subseteq \mathcal{D}(i)\}}$ . We use the simplified notation  $f_T$  instead of  $f_T(\mathcal{D})$  when the meaning is clear. Note that we may view a row  $\mathcal{D}(i)$  of  $\mathcal{D}$  as a one-row database in its own right; hence,  $f_T(\mathcal{D}(i))$  equals 1 if  $\mathcal{D}(i)$  contains  $T$ , and equals 0 otherwise.

We consider *four* different sketching problems that each capture a natural notion of approximate itemset frequency analysis. All four sketching problems permit randomized sketching algorithms, and require that the sketching algorithm succeeds with high probability. Success can be interpreted in two different ways: (1) with high probability, for all  $k$ -itemsets the answer is correct; or (2) for each  $k$ -itemset, with high probability the answer is correct (but it may be very unlikely that one can recover accurate estimates for all  $k$ -itemsets from the sketch simultaneously). These two different notions of success are often termed, respectively, “for all” and “for each” in the compressed sensing literature [AKW14]. Our first two problem definitions (Definitions 1 and 2) correspond to the “for all” notion, while our latter two problem definitions (Definitions 3 and 4) correspond to the weaker “for each” notion.

**Definition 1** (For-All-Itemset-Frequency-Indicator sketches). *A For-All-Itemset-Frequency-Indicator sketch is a tuple  $(\mathcal{S}, \mathcal{Q})$ . The first term  $\mathcal{S}$  is a randomized sketching algorithm. It takes as input a database  $\mathcal{D} \in (\{0, 1\}^d)^n$ , a precision  $\epsilon$ , an itemset size  $k$ , and a failure probability  $\delta$ . It outputs a summary  $\mathcal{S}(\mathcal{D}, k, \epsilon, \delta) \in \{0, 1\}^s$  where  $s$  is the size of the sketch in bits. The second term is a deterministic query procedure  $\mathcal{Q} : \{0, 1\}^s \times \{0, 1\}^d \rightarrow \{0, 1\}$ . It takes as input a summary  $\mathcal{S}$  and a  $k$ -itemset  $T$  and outputs a single bit indicating whether  $T$  is frequent in  $\mathcal{D}$  or not. More precisely, for a triple of input parameters  $(k, \epsilon, \delta)$ , the following two conditions must hold with probability  $1 - \delta$  over the randomness of the sketching algorithm  $\mathcal{S}$ , for every database  $\mathcal{D}$ :*

$$\forall k\text{-itemsets } T \text{ s.t. } f_T > \epsilon, \quad \mathcal{Q}(\mathcal{S}(\mathcal{D}, k, \epsilon, \delta), T) = 1, \text{ and} \quad (1)$$

$$\forall k\text{-itemsets } T \text{ s.t. } f_T < \epsilon/2, \quad \mathcal{Q}(\mathcal{S}(\mathcal{D}, k, \epsilon, \delta), T) = 0. \quad (2)$$

*Note that if  $\epsilon/2 \leq f_T \leq \epsilon$  then either bit value can be returned.*

**Definition 2** (For-All-Itemset-Frequency-Estimator sketches). *A For-All-Itemset-Frequency-Estimator sketch is a tuple  $(\mathcal{S}, \mathcal{Q})$ . Here  $\mathcal{S}$  is defined as above but  $\mathcal{Q} : \{0, 1\}^s \times \{0, 1\}^d \rightarrow [0, 1]$  outputs an approximate frequency. To be precise, the pair  $(\mathcal{S}, \mathcal{Q})$  is a valid For-All-Itemset-Frequency-Estimator sketch for a triple of input parameters  $(k, \epsilon, \delta)$  if for every database  $\mathcal{D}$ :*

$$\Pr[\forall k\text{-itemsets } T, \quad |\mathcal{Q}(\mathcal{S}(\mathcal{D}, k, \epsilon, \delta), T) - f_T| \leq \epsilon] \geq 1 - \delta. \quad (3)$$

**Definition 3** (For-Each-Itemset-Frequency-Indicator sketches). *A For-Each-Itemset-Frequency-Indicator sketch is identical to a For-All-Itemset-Frequency-Indicator sketch, except that Equations (1) and (2) are replaced with the requirement that for every database  $\mathcal{D}$  and any (single)  $k$ -itemset  $T$ :*

*If  $f_T > \epsilon$ , then  $\mathcal{Q}(\mathcal{S}(\mathcal{D}, k, \epsilon, \delta), T) = 1$  with probability at least  $1 - \delta$ , and*

*If  $f_T < \epsilon/2$ , then  $\mathcal{Q}(\mathcal{S}(\mathcal{D}, k, \epsilon, \delta), T) = 0$  with probability at least  $1 - \delta$ .*

**Definition 4** (For-Each-Itemset-Frequency-Estimator sketches). A *For-Each-Itemset-Frequency-Estimator sketch* is identical to a *For-All-Itemset-Frequency-Estimator sketch*, except that Equation (3) is replaced with the requirement that for every database  $\mathcal{D}$  and any (single)  $k$ -itemset  $T$ :  $\Pr[|\mathcal{Q}(\mathcal{S}(\mathcal{D}, k, \epsilon, \delta), T) - f_T| \leq \epsilon] \geq 1 - \delta$ .

**Definition 5.** The space complexity of a sketch, denoted by  $|\mathcal{S}(n, d, k, \epsilon, \delta)|$ , is the maximum sketch size generated by  $\mathcal{S}$  for any database with  $n$  rows and  $d$  columns. That is,  $|\mathcal{S}(n, d, k, \epsilon, \delta)| = \max_{\mathcal{D} \in (\{0,1\}^d)^n} |\mathcal{S}(\mathcal{D}, k, \epsilon, \delta)|$ .

For brevity, we typically omit the parameters  $(n, d, k, \epsilon, \delta)$  when the meaning is clear, and simply write  $|\mathcal{S}|$  to denote the space complexity of a sketch.

## 1.4 Techniques

At a high level our lower bounds are proven via the standard approach of encoding arguments. That is, to prove that  $\Omega(s)$  bit sketches are necessary to solve one of the problems above, we construct a distribution over  $s$ -bit databases  $\mathcal{D} \in (\{0, 1\}^d)^{s/d}$  such that 1) the entropy of the distribution is  $\Omega(s)$  and 2) any itemset frequency sketch can be used to reconstruct the database. Thus, the sketch must have size  $\Omega(s)$  bits in the worst case.

For the simplest case of  $k = 1$  and  $\epsilon = 1/3$ , it is easy to show that  $\Omega(d)$  bits are necessary to solve every version of the itemset sketching problem, since any non-trivial estimation of the 1-itemset frequencies is sufficient to exactly encode a database consisting of a single row. In order to prove larger space lower bounds for larger values of  $k$  and smaller values of  $\epsilon$ , we must show that if we are given either  $k$ -itemset queries for  $k = \omega(1)$ , or sketches with accuracy  $\epsilon = o(1)$ , then we can encode databases consisting of more than one row.

To do so, we combine some new arguments with information-theoretic techniques that were previously developed to solve problems in *differential privacy* [DMNS06]. The problem of constructing differentially private sketches for frequent itemset queries has been studied intensely in recent years (see e.g. [BCD<sup>+</sup>07, KRSU10, De12, GHRU13, TUV12, CTUW14, BUV14]; in these works frequent itemset queries are called monotone conjunction queries and sometimes contingency tables). It turns out that the techniques in these works provide exactly the information-theoretic tools that we need to devise our encoding arguments. Although the connection between the two problems appears coincidental, we believe that information-theoretic tools from the privacy literature may find future applications outside of privacy.<sup>3</sup>

We now sketch roughly how our encoding arguments work. Suppose we have already proven a lower bound of  $\Omega(s)$  bits for any sufficiently accurate  $k$ -itemset sketch via an encoding argument. For this informal discussion, the precise definition of accuracy will not be important. We then have a high-entropy distribution  $\mathcal{D}$  on databases  $D \in (\{0, 1\}^d)^{s/d}$  such that a sufficiently accurate sketch for  $k$ -itemset frequency queries on  $D$  must encode  $D$ . For example, as we discussed above, we can trivially start with a lower bound of  $d$  bits for any non-trivial approximation to the 1-itemset queries, although sometimes we will need to start with

---

<sup>3</sup>Indeed, our use of techniques from the privacy literature may not be purely coincidental, as there is a formal (though quantitatively loose) connection between the problems of developing private and non-private sketches for itemset frequency queries. Suppose that there is a way to take any dataset  $D \in (\{0, 1\}^d)^n$  and create a sketch  $\mathcal{S}$  of  $s$  bits that encodes the answer to every itemset query  $f_T(D)$  to within additive error  $\pm \epsilon$ . Then we can obtain a differentially private sketch that encodes the answers to every query  $f_T(D)$  to within additive error  $\epsilon + O(s/n)$  as follows. Output a sketch  $\mathcal{S}$  with probability proportional to  $\exp(-n \cdot \max_T |f_T(D) - \mathcal{Q}(\mathcal{S}, T)|)$ . Our claims that the output  $\mathcal{S}$  will be differentially private and that with high probability the chosen sketch will have additive error  $\epsilon + O(s/n)$  can both be proven by an elementary analysis, or by using the fact that this algorithm is a special case of the exponential mechanism [MT07] and thus standard results can be applied. Now, suppose that we had a lower bound saying that any differentially private algorithm must incur error  $t/n$  (lower bounds in differential privacy are often of the form  $t/n$  for some  $t$  independent of  $n$ ). Then we would immediately obtain a lower bound saying that any  $\epsilon$ -accurate sketch for itemsets requires  $s = \Omega(t - \epsilon n)$  bits. Thus, accuracy lower bounds in differential privacy *immediately* imply space lower bounds for the associated sketching problem. However, this generic connection rarely leads to bounds that are quantitatively tight.



stronger lower bounds. We then obtain a lower bound of  $\Omega(t \cdot s)$  for any sufficiently accurate sketch for  $k'$ -itemset frequency queries using the following amplification technique, inspired by the technique in [BUV14] for amplifying lower bounds in differential privacy. Roughly, the technique allows us to construct a new distribution  $\mathcal{D}'$  on databases  $D' \in (\{0, 1\}^{2d})^{ts/d}$  such that any sufficiently accurate sketch for  $k'$ -itemset frequency queries can be used to reconstruct an accurate  $k$ -itemset sketch on each of  $t$  different draws from the distribution  $\mathcal{D}$  over  $(\{0, 1\}^d)^{s/d}$ . Since we have started by assuming that such a sketch for  $k$ -itemset frequency queries requires  $\Omega(s)$  bits, we conclude that any sufficiently accurate sketch of for  $k'$ -itemset frequency queries requires  $\Omega(t \cdot s)$  bits of space. The above outline is overly simplified, and the resulting reconstruction will only be approximate. Thus, we need to make sure that the approximation is good enough for our arguments to go through, especially in some of our bounds that require applying the above amplification arguments twice, where the approximation becomes even weaker.

## 2 Naïve upper bounds

In the following we describe three trivial sketching algorithms.

**Definition 6** (RELEASE-DB). *This algorithm simply releases the database verbatim. In other words, the function  $\mathcal{S}$  is the identity and  $\mathcal{Q}$  is a standard database query.*

The space complexity of RELEASE-DB is clearly  $|\mathcal{S}| = O(nd)$  and it produces exact estimates for both For-All-Itemset-Frequency-Estimator and For-All-Itemset-Frequency-Indicator sketches and their For-Each analogs.

**Definition 7** (RELEASE-ANSWERS). *This algorithm precomputes and stores the results to all possible queries.*

Since there are  $\binom{d}{k}$  possible  $k$ -itemset queries, the space complexity of RELEASE-ANSWERS is  $|\mathcal{S}| = O(\binom{d}{k})$  for For-All-Itemset-Frequency-Indicator sketches and their For-Each analogs, and  $|\mathcal{S}| = O\left(\binom{d}{k} \log(1/\epsilon)\right)$  for For-All-Itemset-Frequency-Estimator sketches and their For-Each analogs. The extra  $\log(1/\epsilon)$  factor is needed to represent frequencies as floating point numbers up to precision  $\epsilon$ .

**Definition 8** (SUBSAMPLE). *This algorithm samples rows uniformly at random with replacement from the database. The samples constitute the sketch  $\mathcal{S}(\mathcal{D}, k, \epsilon, \delta)$ . The recovery algorithm  $\mathcal{Q}(\mathcal{S}(\mathcal{D}), T)$  returns the frequency of  $T$  in the sampled rows via a standard database query.*

**Lemma 9** (Subsampling). *SUBSAMPLE outputs*

- *a valid For-All-Itemset-Frequency-Indicator sketch of space complexity  $|\mathcal{S}| = O\left(\epsilon^{-1} d \log\left(\binom{d}{k}/\delta\right)\right)$ ,*
- *a valid For-All-Itemset-Frequency-Estimator sketch of space complexity  $|\mathcal{S}| = O\left(\epsilon^{-2} d \log\left(\binom{d}{k}/\delta\right)\right)$ ,*
- *a valid For-Each-Itemset-Frequency-Indicator sketch of space complexity  $|\mathcal{S}| = O\left(\epsilon^{-1} d \log(1/\delta)\right)$ ,  
and*
- *a valid For-Each-Itemset-Frequency-Estimator sketch of space complexity  $|\mathcal{S}| = O\left(\epsilon^{-2} d \log(1/\delta)\right)$ .*

*Proof.* Each row sample requires  $d$  bits to represent. Thus to prove each of the above statements it suffices to bound the number of row samples required to ensure the accuracy goal. We can do so using standard probabilistic inequalities: Chernoff bounds for sums of independent random variables and union bounds. We will need the following standard forms of the Chernoff bound.

**Lemma 10.** Suppose  $X_1, \dots, X_s$  are independent random  $\{0, 1\}$ -valued random variables with expectation  $p$ , and let  $\bar{X} = \frac{1}{s} \sum_{i=1}^s X_i$ . Then for any  $\epsilon < 2e - 1$ ,  $\mathbb{P}[\bar{X} \notin [(1 - \epsilon)p, (1 + \epsilon)p]] \leq 2 \exp(-sp\epsilon^2/4)$ .

**Lemma 11.** Suppose  $X_1, \dots, X_s$  are independent random  $\{0, 1\}$ -valued random variables with expectation  $p$ , and let  $\bar{X} = \frac{1}{s} \sum_{i=1}^s X_i$ . Then for any  $\epsilon < 1$ ,  $\mathbb{P}[\bar{X} \notin [p - \epsilon, p + \epsilon]] \leq 2 \exp(-2s\epsilon^2)$ .

*For-Each-Itemset-Frequency-Indicator sketches:* Fix a dataset  $\mathcal{D}$  and an itemset  $T$  and let  $p = f_T(\mathcal{D})$ . Consider drawing  $s$  independent uniform random samples of rows  $\mathcal{D}'(1), \dots, \mathcal{D}'(s)$  with replacement from  $\mathcal{D}$ . Let  $\mathcal{D}'$  be the database consisting of the  $s$  row samples. For  $i = 1, \dots, s$ , define the random variable  $X_i = 1$  if  $T \subseteq \mathcal{D}'(i)$  and 0 otherwise. Let  $\bar{X} = f_T(\mathcal{D}') = \frac{1}{s} \sum_{i=1}^s X_i$ . Since the samples  $\mathcal{D}'(i)$  are independent, the random variables  $X_i$  are independent. Moreover, for every  $i$ ,  $\mathbb{E}[X_i] = p$ , and by linearity of expectation  $\mathbb{E}[\bar{X}] = p$ . Then by Lemma 10, we have that

$$\mathbb{P}[f_T(\mathcal{D}') \notin [p/2, 2p]] \leq 2 \exp(-sp/16).$$

The right hand side will be at most  $\delta$  if  $s \geq 16 \ln(2/\delta)/p$  for a sufficiently large constant  $C$ . From this bound, we can deduce that, for For-Each-Itemset-Frequency-Indicator sketches, it suffices to take  $s = O(\epsilon^{-1} \log(1/\delta))$  row samples to ensure that the accuracy requirement of Definition 3 is satisfied.

*For-Each-Itemset-Frequency-Estimator sketches:* The setup is the same, except that we apply Lemma 11 instead of Lemma 10 to obtain:

$$\mathbb{P}[f_T(\mathcal{D}') \notin [p - \epsilon, p + \epsilon]] \leq 2 \exp(-2s\epsilon^2).$$

The right hand side will be at most  $\delta$  if  $s \geq \epsilon^{-2} \ln(2/\delta)$ . Thus, for For-Each-Itemset-Frequency-Estimator sketches it suffices to take  $s = O(\epsilon^{-2} \log(1/\delta))$  row samples to ensure that the accuracy requirement of Definition 4 is satisfied.

*For-All-Itemset-Frequency-Indicator sketches:* By the analysis above, we know that for any  $\delta' > 0$  and any itemset  $T$ ,  $f_T(\mathcal{D}')$  is accurate with probability at least  $1 - \delta'$  if  $s = O(\epsilon^{-1} \log(1/\delta'))$ . Thus, by a union bound

$$\mathbb{P}[\exists T \subseteq [d], |T| = k, f_T(\mathcal{D}') \text{ is not accurate}] \leq \binom{d}{k} \mathbb{P}[f_T(\mathcal{D}') \text{ is not accurate}] \leq \binom{d}{k} \delta'.$$

Now, by setting  $\delta' = \delta / \binom{d}{k}$ , we can see that it suffices to take  $s = O(\epsilon^{-1} \log(\binom{d}{k}/\delta))$  samples to ensure accuracy.

*For-All-Itemset-Frequency-Estimator sketches:* Here we apply the same union bound argument to our analysis of For-Each-Itemset-Frequency-Estimator sketches. We can easily see that it suffices to take  $s = O(\epsilon^{-2} \log(\binom{d}{k}/\delta))$ .  $\square$

For any setting of the parameters  $(n, d, k, \epsilon, \delta)$ , the minimal space usage among the above three trivial algorithms constitutes our naïve upper bound for all four sketching problems that we consider, formalized in Theorem 12 below.

**Theorem 12.** (a) For any  $(k, \epsilon, \delta)$ , there is a randomized algorithm that, given any database  $\mathcal{D} \in (\{0, 1\}^d)^n$ , outputs a For-All-Itemset-Frequency-Indicator sketch of size  $O\left(\min\left\{nd, \binom{d}{k}, \epsilon^{-1}d \log\left(\binom{d}{k}/\delta\right)\right\}\right)$ .

(b) For any  $(k, \epsilon, \delta)$ , there is a randomized algorithm that, given any database  $\mathcal{D} \in (\{0, 1\}^d)^n$ , outputs a For-All-Itemset-Frequency-Estimator sketch of size  $O\left(\min\left\{nd, \binom{d}{k} \log(1/\epsilon), \epsilon^{-2}d \log\left(\binom{d}{k}/\delta\right)\right\}\right)$ .

(c) For any  $(k, \epsilon, \delta)$ , there is a randomized algorithm that, given any database  $\mathcal{D} \in (\{0, 1\}^d)^n$ , outputs a For-Each-Itemset-Frequency-Indicator sketch of size  $O\left(\min\left\{nd, \binom{d}{k}, \epsilon^{-1}d \log(1/\delta)\right\}\right)$ .

(d) For any  $(k, \epsilon, \delta)$ , there is a randomized algorithm that, given any database  $\mathcal{D} \in (\{0, 1\}^d)^n$ , outputs a For-Each-Itemset-Frequency-Estimator sketch of size  $O\left(\min\left\{nd, \binom{d}{k} \log(1/\epsilon), \epsilon^{-2}d \log(1/\delta)\right\}\right)$ .

### 3 Lower Bounds

In this section, we turn to proving lower bounds on the size of For-All-Itemset-Frequency-Indicator and For-All-Itemset-Frequency-Estimator sketches. Notice that the algorithms RELEASE-ANSWERS and SUBSAMPLE produce sketches whose size is independent of  $n$ ; hence, it is impossible to prove lower bounds that grow with  $n$ . Consequently, we state our lower bounds in terms of the parameters  $(d, k, 1/\epsilon)$ , with all of our lower bounds holding as long as  $n$  is sufficiently large relative to these three parameters. This parameter regime — with  $n$  a sufficiently large polynomial in  $d$ ,  $k$ , and  $1/\epsilon$  — is consistent with typical usage scenarios, where the number of rows in a database far exceeds the number of attributes. In our formal theorem statements, we make explicit precisely how large a polynomial  $n$  must be in terms of  $d$ ,  $k$ , and  $1/\epsilon$  for the lower bound to hold.

Each of our lower bounds also requires  $d$ ,  $k$ , and  $1/\epsilon$  to satisfy certain mild technical relationships with each other — for example, Theorems 13 and 14 require that  $1/\epsilon < \binom{d/2}{k-1}$ . In all cases, the assumed technical relationship between the parameters is necessary or close to necessary for the claimed lower bound to hold. For instance, the  $\Omega(d/\epsilon)$  lower bound of Theorems 13 and 14 is false for  $1/\epsilon \gg \binom{d/2}{k-1}$ , as the algorithm RELEASE-ANSWERS would output a sketch of size  $o(d/\epsilon)$  in this parameter regime.

#### 3.1 Overview of the Lower Bounds

We now formally state all of the lower bounds that we prove, and place our results in context. Throughout this section, we assume that the failure probability  $\delta$  of the sketching algorithm is a constant less than one.

**Resolving the complexity of For-Each-Itemset-Frequency-Indicator sketches.** The main result stated in this section is a relatively simple  $\Omega(d/\epsilon)$  lower bound for the For-Each-Itemset-Frequency-Indicator sketching problem (Theorem 14 below), which is essentially optimal despite its simplicity. For expository purposes, it will be convenient to first state an analogous lower bound for the (harder) For-All-Itemset-Frequency-Indicator sketching problem. In Section 3.2, we prove the For-All lower bound first, and then explain how to modify the proof to handle the For-Each case.

**Theorem 13.** *Let  $k \geq 2$ . Suppose that  $1/\epsilon \leq \binom{d/2}{k-1}$ , and the failure probability  $\delta < 1$  is constant. Then for  $n \geq 1/\epsilon$ , the space complexity of any valid For-All-Itemset-Frequency-Indicator sketch is  $|\mathcal{S}(n, k, d, \epsilon, \delta)| = \Omega(d/\epsilon)$ .*

**Theorem 14.** *Let  $k \geq 2$ . Suppose that  $1/\epsilon \leq \binom{d/2}{k-1}$ , and  $\delta < 1/3$ . Then for  $n \geq 1/\epsilon$ , the space complexity of any valid For-Each-Itemset-Frequency-Indicator sketch is  $|\mathcal{S}(n, k, d, \epsilon, \delta)| = \Omega(d/\epsilon)$ .*

Theorem 14 is tight whenever it applies (i.e., when  $1/\epsilon < \binom{d/2}{k-1}$ ), as it matches the  $O(d/\epsilon)$  upper bound obtained by the algorithm SUBSAMPLE for the For-Each-Itemset-Frequency-Indicator sketching problem (see Theorem 12). And the algorithm RELEASE-ANSWERS achieves a summary size of  $\binom{d}{k}$  for the For-Each-Itemset-Frequency-Indicator sketching problem, which is asymptotically optimal when  $1/\epsilon \geq \binom{d/2}{k-1}$  and  $k = O(1)$ . Therefore, Theorems 12 and 14 together precisely resolve the complexity of For-Each-Itemset-Frequency-Indicator sketches for all values of  $d$  and  $\epsilon$ , when  $k = O(1)$ .



**Resolving the complexity of For-All-Itemset-Frequency-Indicator sketches.** Theorem 13 is tight for For-All-Itemset-Frequency-Indicator sketches when  $1/\epsilon$  is small relative to the other input parameters  $n$  or  $d$ . In particular, when  $n = 1/\epsilon$ , RELEASE-DB provides a trivial matching sketch that is  $O(nd) = O(d/\epsilon)$  bits in size. In addition, when  $k = O(1)$  and  $1/\epsilon \geq \binom{d/2}{k-1}$ , RELEASE-ANSWERS provides a matching sketch that is  $O(\binom{d}{k}) = O(d/\epsilon)$  bits in size. The tightness of Theorem 13 in these parameter regimes is arguably surprising, as it shows that the For-All-Itemset-Frequency-Indicator sketching problem is *equivalent* in complexity to its For-Each analog in these regimes.

However, when  $1/\epsilon \ll \binom{d/2}{k-1}$ , Theorem 13 is not tight for For-All-Itemset-Frequency-Indicator sketches, because it has suboptimal dependence on  $d$  and  $k$ . Our main result for the For-All-Itemset-Frequency-Indicator sketching problem establishes a tight lower bound, matching the  $O(\epsilon^{-1} d \log \binom{d}{k})$  upper bound for the problem obtained by the algorithm SUBSAMPLE.

**Theorem 15.** *Let  $k \geq 3$ , and suppose that  $1/\epsilon = O\left(\binom{d/3}{\lfloor (k-1)/2 \rfloor}\right)$  and the failure probability  $\delta < 1$  is a constant. Then for any  $n \geq kd \log(d/k)/\epsilon$ , the space complexity of any valid For-All-Itemset-Frequency-Indicator sketch is  $|\mathcal{S}(n, k, d, \epsilon, \delta)| = \Omega(kd \log(d/k)/\epsilon)$ .*

**Comparison to Price’s work.** Price independently proved an  $\Omega(d \log(d)/\epsilon)$  lower bound on the size of For-All-Itemset-Frequency-Indicator sketches for  $k \geq 2$  and  $1/\epsilon \leq d^{.99}$  [Pri14]. This matches the lower bound of Theorem 15 for any  $k = O(1)$ , but not for  $k = \omega(1)$ . One advantage of Price’s result is that it holds for  $k = 2$  and any value of  $\epsilon$  satisfying  $1/\epsilon \leq d^{.99}$ ; in contrast, Theorem 15 holds for  $k \geq 3$ . We remark that our proof of Theorem 15 actually establishes the  $\Omega(d \log(d))$  lower bound even for  $k = 2$ , but our extension of the proof to sub-constant values of  $\epsilon$  requires  $k \geq 3$ .

**Essentially resolving the complexity of For-All-Itemset-Frequency-Estimator sketches.** Theorem 16 below establishes a lower bound for the For-All-Itemset-Frequency-Estimator sketching problem with a quadratically stronger dependence on  $1/\epsilon$ , relative to the linear dependence that is necessary and sufficient for the For-All-Itemset-Frequency-Indicator problem. Our lower bound matches the space usage of SUBSAMPLE (cf. Lemma 9) up to a factor of  $\log_{(q)}(1/\epsilon)$  for any desired constant  $q > 0$ , where  $\log_{(q)}$  denotes the logarithm function iterated  $q$  times (e.g.  $\log_{(3)}(x) = \log \log \log(x)$ ).<sup>4</sup>

**Theorem 16.** *Fix any integer constants  $c \geq 2$ ,  $q \geq 1$ , and let  $\delta < 1$  be a constant. Let  $k \geq c + 1$ , let  $d, \epsilon$  satisfy  $1/\epsilon^2 \leq d^{c-1}/\log_{(q)}(1/\epsilon^2)$ , and let  $\delta < 1$  be a constant. Let  $\mathcal{S}$  be a For-All-Itemset-Frequency-Estimator sketching algorithm capable of answering all  $k$ -itemset frequency queries to error  $\pm \epsilon$  on databases  $\mathcal{D} \in (\{0, 1\}^d)^n$  for any  $n > d \log(d/k) \log_{(q)}(1/\epsilon^2)/\epsilon^2$ . Then  $|\mathcal{S}(n, d, k, \epsilon, \delta)| = \Omega\left(\frac{kd \log(d/k)}{\epsilon^2 \log_{(q)}(1/\epsilon)}\right)$ .*

For illustration, consider fixing the constants  $c = q = 10$ . Then theorem says that for every  $k \geq 11$ , if  $1/\epsilon = O(d^{4.5}/\log_{(10)}(d))$ , and the database size  $n$  is sufficiently large, then there is a space lower bound of  $|\mathcal{S}| = \Omega\left(kd \log(d/k)/\left(\epsilon^2 \log_{(10)}(1/\epsilon)\right)\right)$ .

<sup>4</sup>An anonymous reviewer has pointed out that subsequent work of Van Gucht et al. [GWZ15, Theorem 6] can be combined with our amplification techniques to remove the  $\log_{(q)}(1/\epsilon)$  factor in the bound of Theorem 16, when  $\epsilon > 1/\sqrt{d}$ ; the resulting bound matches the size of the sketch produced by SUBSAMPLE up to a constant factor. (Note that Theorem 16 holds even for values of  $\epsilon$  smaller than  $1/\sqrt{d}$ , whenever  $k > 3$ .) In more detail, the proof of [GWZ15, Theorem 6] implies an  $\Omega(d/\epsilon^2)$  lower bound on the size of any For-All-Itemset-Frequency-Estimator sketch answering 99% of all 2-itemset queries, when  $\epsilon > 1/\sqrt{d}$ . By combining this result with our amplification techniques, it is possible to prove an  $\Omega\left(\frac{kd \log(d/k)}{\epsilon^2}\right)$  lower bound on the size of any For-All-Itemset-Frequency-Estimator sketch answering all  $k$ -itemset queries for  $k \geq 3$ , whenever  $\epsilon > 1/\sqrt{d}$ .

Finally, we use a simple reduction to show that the above theorem also implies a lower bound for the For-Each case, which is optimal up to a factor of  $\log_{(q)}(1/\epsilon)$ .<sup>5</sup>

**Theorem 17.** *Fix any integer constants  $c \geq 2$ ,  $q \geq 1$ . Let  $k \geq \max\{3, c + 1\}$ , let  $d, \epsilon$  satisfy  $1/\epsilon^2 \leq d^{c-1}/\log_{(q)}(1/\epsilon^2)$ , and let  $\delta < 1/2$  be a constant. Let  $\mathcal{S}$  be an For-Each-Itemset-Frequency-Estimator sketching algorithm capable of answering any (single)  $k$ -itemset frequency queries to error  $\pm\epsilon$  on databases  $\mathcal{D} \in (\{0, 1\}^d)^n$  for any  $n > d \log(d/v) \log_{(q)}(1/\epsilon^2)/\epsilon^2$ . Then  $|\mathcal{S}(n, d, k, \epsilon, \delta)| = \Omega\left(\frac{d}{\epsilon^2 \log_{(q)}(1/\epsilon)}\right)$ .*

## 3.2 Lower Bound Proofs for Itemset-Frequency-Indicator Sketches

### 3.2.1 First Lower Bound Proofs: Theorems 13 and 14

We begin by proving our two simplest bounds (Theorems 13 and 14). Recall that the former applies to For-All-Itemset-Frequency-Indicator sketches, and the latter applies even to their For-Each analogs. The proofs consider databases in which even a single appearance of an itemset already makes it frequent. We show that, unsurprisingly, essentially no compression is possible in this setting. (For simplicity, we assume that  $1/\epsilon$  is an integer throughout.)

*Proof of Theorem 13.* Our proof uses an encoding argument. Consider the following family of databases. There will be  $1/\epsilon$  possible settings for each row; as  $n \geq 1/\epsilon$ , some rows may be duplicated. For expository purposes, we begin by describing the setting with  $n = 1/\epsilon$ , in which case there are no duplicated rows. The first  $d/2$  columns in each row contain a unique set of exactly  $k - 1$  attributes. The last  $d/2$  attributes in each row are unconstrained. The only minor technicality is that to ensure that each row can receive a unique set of  $k - 1$  items from the first  $d/2$  attributes, we require  $1/\epsilon \leq \binom{d/2}{k-1}$ .

Given a valid For-All-Itemset-Frequency-Indicator or For-All-Itemset-Frequency-Estimator sketch for this database, one can recover all of the values  $\mathcal{D}(i, j)$  where  $j \geq d/2$  as follows. For any  $j \geq d/2$ , let  $T_{i,j}$  be a set of  $k$  attributes, where the first  $k - 1$  attributes in  $T_{i,j}$  correspond to the  $k - 1$  attributes in the first  $d/2$  columns in the  $i$ th row, and the final attribute in  $T_{i,j}$  is  $j$ . Notice that  $T_{i,j} \in \mathcal{D}$  if and only if  $\mathcal{D}(i, j) = 1$ . Moreover, since  $n = 1/\epsilon$  we have that  $f_T \geq \epsilon$  if and only if  $\mathcal{D}(i, j) = 1$ . Given a valid For-All-Itemset-Frequency-Indicator or For-All-Itemset-Frequency-Estimator sketch for this database, one can iterate over all itemsets  $T_{i,j}$  to recover all the values  $\mathcal{D}(i, j)$  where  $j \geq d/2$ . Since these are an unconstrained set of  $d/(2\epsilon)$  bits, the space complexity of storing them (with  $1 - \Omega(1)$  failure probability) is  $\Omega(d/\epsilon)$  by standard information theory.

For  $n$  a multiple of  $1/\epsilon$ , we construct a database with  $1/\epsilon$  rows as above, and duplicate each row  $n\epsilon$  times; in this case we have  $f_T \geq \epsilon$  if and only if  $\mathcal{D}(i, j) = n\epsilon$ . More generally, when  $n \geq 1/\epsilon$ , duplicating each row at least  $\lfloor n\epsilon \rfloor$  times, we have  $f_T \geq \epsilon$  if and only if  $\mathcal{D}(i, j) \geq \lfloor n\epsilon \rfloor$ .  $\square$

We remark that the condition  $1/\epsilon \leq \binom{d/2}{k-1}$  can be relaxed to  $1/\epsilon \leq \binom{\alpha d}{k-1}$  for any constant  $\alpha < 1$ , by a simple modification of the proof. We now extend the argument used to prove Theorem 13 to the For-Each case.

*Proof of Theorem 14.* Recall that in the setting of one-way randomized communication complexity, there are two parties, Alice and Bob. Alice has an input  $x \in \mathcal{X}$ , Bob has an input  $y \in \mathcal{Y}$ , and Alice and Bob both have access to a public random string  $r$ . Their goal is to compute  $f(x, y)$  for some agreed upon function  $f: \mathcal{X} \times \mathcal{Y} \rightarrow \{0, 1\}$ . Alice sends a single message  $m(x, r)$  to Bob. Based on this message, Bob outputs a bit, which is required to equal  $f(x, y)$  with probability at least  $2/3$ .

---

<sup>5</sup>Footnote 4 implies that one can remove the  $\log_{(q)}(1/\epsilon)$  factor in the bound of Theorem 17 when  $\epsilon > 1/\sqrt{d}$ . The resulting bound matches the size of the sketch produced by SUBSAMPLE in this parameter regime up to a constant factor.

We consider the well-known INDEX function. In this setting, Alice's input  $x$  is an  $N$ -dimensional binary vector, Bob's input  $y$  is an index in  $[N]$ , and  $f(x, y) = x_y$ , the  $y$ 'th bit of  $x$ . It is well-known that one-way randomized communication protocols for INDEX require communication  $\Omega(N)$  [Abl96]. We show how to use any For-Each-Itemset-Frequency-Indicator sketching algorithm  $\mathcal{S}$  to obtain a one-way communication protocol for INDEX on vectors of length  $N = (d/2) \cdot 1/\epsilon$ , with communication proportional to  $|\mathcal{S}(n, d, \epsilon, k, \delta)|$ .

Specifically, let  $(n, d, k, \epsilon, \delta)$  be as in the statement of the theorem. Consider any Boolean vector  $x \in \{0, 1\}^N$ , where  $N = (d/2) \cdot 1/\epsilon$ . We associate each index  $y \in [N]$  with a unique  $k$ -itemset  $T_y \subseteq [d]$  of the following form: the first  $k - 1$  attributes in  $T_y$  are each in  $[d/2]$ , and the final attribute in  $T_y$  is in  $\{d/2 + 1, \dots, d\}$ . The proof of Theorem 13 established the following fact: given any vector  $x \in \{0, 1\}^N$ , there exists a database  $\mathcal{D}_x$  with  $d$  columns and  $n$  rows satisfying the following two properties for all  $y \in [N]$ :

$$x_y = 1 \implies f_{T_y}(\mathcal{D}_x) \geq \epsilon, \text{ and } x_y = 0 \implies f_{T_y}(\mathcal{D}_x) = 0 < \epsilon/2. \quad (4)$$

Hence, we obtain a one-way randomized protocol for the INDEX function as follows: Alice sends to Bob  $\mathcal{S}(\mathcal{D}_x, k, \epsilon, \delta)$  at a total communication of  $|\mathcal{S}(n, d, \epsilon, k, \delta)|$  bits, and Bob outputs  $\mathcal{Q}(\mathcal{S}(\mathcal{D}_x, k, \epsilon, \delta), T_y)$ . It follows immediately from Equation (4) and Definition 3 that Bob's output equals  $x_y$  with probability  $1 - \delta$ . We conclude that  $|\mathcal{S}(n, d, \epsilon, k, \delta)| = \Omega(N) = \Omega(d/\epsilon)$ , completing the proof.  $\square$

### 3.2.2 A Tight Lower Bound for For-All-Itemset-Frequency-Indicator Sketches: Proof of Theorem 15

Our proof of Theorem 15 is inspired by an approach from [BUV14] for “bootstrapping” two weak privacy lower bounds into a stronger lower bound.

*Proof of Theorem 15.* We begin by proving an  $\Omega(kd \log(d/k))$  lower bound for  $\epsilon = 1/50$  and every  $k \geq 2$  (the specific choice  $\epsilon = 1/50$  is for convenience; the lower bound holds for any suitably small constant). We then use the ideas underlying the proof of Theorem 14 to extend the lower bound to sub-constant values of  $\epsilon$ , for any  $k \geq 3$ .

**The case of  $\epsilon = 1/50$ .** We begin by recalling a basic combinatorial fact about  $k$ -itemset frequency queries on databases with  $d$  attributes.

**Fact 18.** *For any  $k' \geq 1$ , let  $v = k' \cdot \log(d/k')$ . There exist strings  $x_1, \dots, x_v \in \{0, 1\}^d$  such that for every string  $s \in \{0, 1\}^v$ , there is a  $k'$ -itemset  $T_s$  such that  $f_{T_s}(x_i) = s_i$  for all  $i \in [v]$ .*

*Proof of Fact 18.* It is well-known that the set of  $k'$ -itemset frequency queries (equivalently,  $k'$ -way monotone conjunction queries), when evaluated on  $d$ -bit vectors (i.e., on  $d$ -attribute database rows), has VC dimension at least  $k' \cdot \log(d/k')$ . The desired strings  $x_1, \dots, x_v$  are simply the shattered set whose existence is guaranteed by having VC dimension  $v$ . It is also not difficult to directly construct the shattered set; we provide such a direct construction in Appendix A for completeness.  $\square$

Let  $k' = k - 1$ , let  $v = (k - 1) \log(d/(k - 1))$ , and let  $x_1, \dots, x_v \in \{0, 1\}^d$  be the strings promised by Fact 18. Let  $y_1, \dots, y_v \in \{0, 1\}^d$  be an arbitrary set of  $v$  strings of length  $d$ . We will show how to construct a database  $\mathcal{D}$  with  $v$  rows and  $2d$  columns such that, with probability at least  $1 - \delta$ , at least 96% of the  $dv$  bits in  $(y_1, \dots, y_v)$  can be reconstructed from any For-All-Itemset-Frequency-Indicator sketch  $\mathcal{S}(\mathcal{D}, k, 1/50, \delta)$ . This will imply an  $\Omega(dv)$  space lower bound. Specifically, define row  $i$  of  $\mathcal{D}$  to be

$$\mathcal{D}(i) := (x_i, y_i). \quad (5)$$

That is, the first  $d$  bits in  $\mathcal{D}(i)$  are equal to  $x_i$  and the last  $d$  bits are equal to  $y_i$ .

The key observation behind the reconstruction of the  $y_i$ 's is that, given *exact* answers to all  $k$ -way itemset frequency queries, one can compute the inner product between the last  $d$  columns of  $\mathcal{D}$  and any desired vector. Moreover, it is easy to see that, given sufficiently many inner products, any column of  $\mathcal{D}$  can be exactly reconstructed. However, For-All-Itemset-Frequency-Indicator sketches do not provide exact answers to itemset frequency queries; they merely indicate whether the frequency of an itemset is larger than  $\epsilon$  or smaller than  $\epsilon/2$ . Nonetheless, we show that in order to reconstruct 96% of the bits in any given column of  $\mathcal{D}$ , it is enough to know, for sufficiently many vectors, whether the inner product of the column with the vector is larger than  $\epsilon = 1/50$  or smaller than  $\epsilon/2 = 1/100$ .

In more detail, fix a column  $j \in \{d+1, \dots, 2d\}$ . Let  $t := (y_{1,j}, \dots, y_{v,j}) \in \{0, 1\}^v$  be the bits in this column of  $\mathcal{D}$ . Fix a string  $s$  in  $\{0, 1\}^v$ , and for any  $j \in [d]$ , let  $T_{s,j} = T_s \cup \{j\}$ , where  $T_s$  is defined as in Fact 18. We claim that the correct answer to the itemset frequency query  $f_{T_{s,j}}(\mathcal{D})$  is  $\langle s, t \rangle / v$ . To see this, notice that by the definition of  $\mathcal{D}$  (Equation (5)),  $T_{s,j}$  is contained in any row  $i$  of  $\mathcal{D}$  such that  $s_i = y_{ij} = 1$ , and  $T_{s,j}$  is not contained in any other rows. Hence,  $f_{T_{s,j}}(\mathcal{D}) = \frac{1}{v} |\{i : s_i = 1, y_{ij} = 1\}| = \langle s, t \rangle / v$ .

It follows that any For-All-Itemset-Frequency-Indicator sketch  $\mathcal{S}$  of  $\mathcal{D}$  that provides answers with error parameter  $\epsilon = 1/50$  for all  $k$ -itemsets, provides a bit  $b_s$  for every  $s \in \{0, 1\}^v$  such that the following holds:  $b_s = 1$  if  $\langle s, t \rangle / v > \epsilon$ , and  $b_s = 0$  if  $\langle s, t \rangle / v < \epsilon/2$ . The following lemma implies that we can use the  $b_s$  values to reconstruct a vector  $t'$  that is close to  $t$  in Hamming distance.

**Lemma 19.** *Suppose for every  $s \in \{0, 1\}^v$ , we are given a bit  $b_s$  satisfying  $b_s = 1$  if  $\langle s, t \rangle / v > \epsilon$ , and  $b_s = 0$  if  $\langle s, t \rangle / v < \epsilon/2$ . Let  $t' \in \{0, 1\}^v$  be any vector that is consistent with all of the  $b_s$  values, in the sense that  $\langle s, t' \rangle / v > \epsilon$  for all  $s$  such that  $b_s = 1$ , and  $\langle s, t' \rangle / v < \epsilon/2$  for all  $s$  such that  $b_s = 0$ . Then the Hamming distance between  $t$  and  $t'$  is at most  $v/25$ .*

*Proof of Lemma 19.* Consider any vector  $t' \in \{0, 1\}^v$  such that  $t$  and  $t'$  differ in more than  $v/25$  fraction of bits. Then there must a set of coordinates  $S \subseteq [v]$  of size at least  $v/50$  such that at least one of the two conditions is satisfied: (a)  $t'_j = 1$  and  $t_j = 0$  for all  $j \in S$ , or (b)  $t'_j = 0$  and  $t_j = 1$  for all  $j \in S$ .

Assume without loss of generality that Condition (a) is satisfied (the proof in the case that Condition (b) is satisfied is analogous). Consider the vector  $s \in \{0, 1\}^v$  that is the indicator vector of  $S$ . Then  $\langle s, t \rangle = 0$ , so  $b_s = 0$ . However,  $\langle s, t' \rangle \geq v/50$ . This implies that  $t'$  is not consistent with the value  $b_s$  returned by the For-All-Itemset-Frequency-Indicator sketch, proving the lemma.  $\square$

Lemma 19 implies that, for any  $k \geq 2$ , given the For-All-Itemset-Frequency-Indicator sketch  $\mathcal{S}(\mathcal{D}, k, 1/50, \delta)$ , we can recover at least 96% of the bits of  $(y_1, \dots, y_v)$  with probability at least  $1 - \delta$ . Suppose we let  $(y_1, \dots, y_v)$  be the error-corrected encoding of a vector  $(y'_1, \dots, y'_z) \in \{0, 1\}^z$ , using a code with constant rate that is uniquely decodable from 4% errors (e.g. using a Justesen code [Jus72]). Then  $z = \Omega(v)$ , and it follows from the above that  $(y'_1, \dots, y'_z)$  can be *exactly* reconstructed from  $\mathcal{S}(\mathcal{D}, k, 1/50, \delta)$  with probability at least  $1 - \delta$ . Hence,  $\mathcal{S}(\mathcal{D})$  allows for exact reconstruction of  $z = \Omega(dv)$  arbitrary bits with probability  $1 - \delta$ . Basic information theory then implies that  $|\mathcal{S}| = \Omega(dv) = \Omega(kd \log(d/k))$ .

**The case of  $\epsilon = o(1)$ .** For any  $k \geq 3$ , suppose that we are given a For-All-Itemset-Frequency-Indicator sketching algorithm  $\mathcal{S}$  that is capable of answering For-All-Itemset-Frequency-Indicator queries with error parameter  $\epsilon$  for all  $k$ -itemsets. For simplicity, we assume  $k$  is odd. At a high level, we show that, given  $m = \frac{1}{50\epsilon}$  independent databases  $\mathcal{D}_i$ , each with  $v$  rows and  $2d$  columns, we can construct a single “larger” database  $\mathcal{D}$  with  $mv$  rows and  $3d$  columns such that the following holds: for every  $\mathcal{D}_i$ ,  $\mathcal{S}(\mathcal{D}, k, \epsilon, \delta)$  can be used to answer all For-All-Itemset-Frequency-Indicator queries on  $\mathcal{D}_i$  with error parameter  $\epsilon' = 1/50$  for all  $(k+1)/2$ -itemsets. Since we have assumed  $k \geq 3$ , it holds that  $(k+1)/2 \geq 2$ ; hence, we can apply our earlier analysis to conclude that any such summary for  $\mathcal{D}_i$  contains  $\Omega(kd \log(d/k))$  bits of information, in the sense that it can encode an arbitrary bit vector of this length. It follows that  $\mathcal{S}(\mathcal{D}, k, \epsilon, \delta)$  contains  $\Omega(kd \log(d/k)/\epsilon)$  bits of information, proving the theorem. Details follow.

Let  $T_1, \dots, T_m \subseteq [d]$  be distinct  $((k-1)/2)$ -itemsets (note that as in the proof of Theorem 13, we can indeed choose  $m$  such  $T_i$ 's as long as  $1/\epsilon < \binom{d}{(k-1)/2}$ ). Now consider any  $m$  independent databases  $\mathcal{D}_1, \dots, \mathcal{D}_m$ , each with  $v$  rows and  $2d$  columns. We construct a new  $(mv) \times 3d$  database  $\mathcal{D}$  by appending the  $d$ -bit indicator vector of  $T_i$  to each row of  $\mathcal{D}_i$ , and letting  $\mathcal{D}$  be the concatenation of all of the resulting databases.

For each  $T_i$ , let  $T' \subseteq [3d]$  be defined via  $T'_i = \{j + 2d : j \in T_i\}$ . That is,  $T'_i$  is simply  $T$  “shifted” to operate on the final  $d$  of the  $3d$  attributes over which the “larger” database  $\mathcal{D}$  is defined. Let  $T^* \subseteq [2d]$  be any  $(k+1)/2$ -itemset, and for each  $i \in [m]$ , let  $T_i^* \subseteq [3d]$  be the  $k$ -itemset defined via:  $T_i^* = T^* \cup T'_i$ . Observe that  $f_{T^*}(\mathcal{D}_i) = m \cdot f_{T_i^*}(\mathcal{D}')$ . Hence,  $f_{T_i^*}(\mathcal{D}) > \epsilon$  if and only if  $f_{T^*}(\mathcal{D}_i) > 1/50$ , and  $f_{T_i^*}(\mathcal{D}) < \epsilon/2$  if and only if  $f_{T^*}(\mathcal{D}_i) < 1/100$ . That is, one can use  $\mathcal{S}(\mathcal{D}, k, \epsilon, \delta)$  to answer all For-All-Itemset-Frequency-Indicator queries on  $\mathcal{D}_i$  with error parameter  $1/50$  for all  $k$ -itemsets (this holds simultaneously for all  $i$  with probability  $1 - \delta$ ).

By the argument for the case  $\epsilon = 1/50$ , this implies that  $\mathcal{S}(\mathcal{D}, k, \epsilon, \delta)$  can be used to losslessly encode an arbitrary vector of length  $\Omega(d \cdot v \cdot m) = \Omega((kd \log(d/k)/\epsilon))$ , and thus  $|\mathcal{S}(\mathcal{D}, k, \epsilon, \delta)| = \Omega((kd \log(d/k)/\epsilon))$ . This completes the proof of Theorem 15.  $\square$

## 4 Lower Bounds Proofs for Itemset-Frequency-Estimator Sketches

### 4.1 The For-All Case: Proof of Theorem 16

#### 4.1.1 Informal Overview of the Proof

For constant  $k' \geq 2$ , an  $\tilde{\Omega}(d/\epsilon^2)$  lower bound on the size of For-All-Itemset-Frequency-Estimator sketches follows fairly directly from existing work in the literature on differential privacy (cf. Kasiviswanathan et al. [KRSU10]; we refer to this work as KRSU). The idea of KRSU’s result is the following. Itemset frequency queries are a linear class of queries, in the sense that we can represent any database as a vector  $z$  (in which each entry of  $z$  corresponds to a possible record in  $\{0, 1\}^d$  and its value is the number of such records in the database), and the vector of answers to all  $k'$ -itemset frequency queries on  $z$  can be written as  $Az$  for some matrix  $A$ . Given a vector  $y$  of approximate answers to these queries, one can try to reconstruct  $z$  via the approximation  $\hat{z} = A^{-1}y$ , where  $A^{-1}$  denotes the Moore-Penrose pseudo-inverse of  $A$  (this is essentially reconstruction via  $L_2$ -distance minimization). If the matrix  $A$  has a “nice” spectrum, then it is possible to bound the distance between  $\hat{z}$  and  $z$ . If this distance is small enough, then any description of  $y$  contains many bits of information, since it essentially encodes an entire database  $z$ .

However, KRSU do not actually look at the matrix  $A$  corresponding to  $k'$ -itemset frequency queries. Instead, they look at a matrix  $M^{(k')}$  they define as follows. Consider a database  $\mathcal{D}$  with  $n$  rows and  $k'$  columns, where the first  $k' - 1$  columns of  $\mathcal{D}$  are generated at random. For any fixed setting of the first  $k' - 1$  columns, the vector of answers to  $k'$ -itemset frequency queries on  $\mathcal{D}$  are a function only of column  $k'$  of  $\mathcal{D}$ . Denoting column  $k'$  of  $\mathcal{D}$  by  $x$ , these answers can be written in the form  $M^{(k')}x$ , for a particular matrix  $M^{(k')}$  derived from the first  $k' - 1$  columns of  $\mathcal{D}$ .

KRSU show that  $M^{(k')}$  behaves a lot like a matrix with truly random entries from  $\{0, 1\}$ , and hence has a “nice” spectrum (with high probability over the random choice of the first  $k' - 1$  columns of  $\mathcal{D}$ ). This ensures that  $\hat{x} = A^{-1}y$  is a “good” approximation to the last column  $x$  of  $\mathcal{D}$  as long as all answers in  $y$  have error  $\epsilon \lesssim \sqrt{n}$ . Put another way, if the error in the answers is  $\epsilon$ , then it is possible to reconstruct column  $k'$  of  $\mathcal{D}$  as long as the number of rows is at most (roughly)  $1/\epsilon^2$ .

This shows that one can use a summary providing  $\epsilon$ -approximate answers to all  $k'$ -itemset frequency queries on a database with  $k'$  columns and  $1/\epsilon^2$  rows to reconstruct  $\tilde{\Omega}(1/\epsilon^2)$  arbitrary bits. It is possible to extend this argument to databases with  $d + k' - 1$  columns and  $\tilde{\Omega}(1/\epsilon^2)$  rows, yielding a  $\tilde{\Omega}(d/\epsilon^2)$  lower bound on the size of For-All-Itemset-Frequency-Estimator sketches for such databases.



Our main contribution for the For-All-Itemset-Frequency-Estimator problem is to combine such an  $\tilde{\Omega}(d/\epsilon^2)$  lower bound for sketches for  $k'$ -way marginals with a technique for “amplifying” the lower bound to  $\tilde{\Omega}(k \cdot \log(d/k) \cdot d/\epsilon^2)$  for  $(k + k')$ -way marginal queries (we used essentially the same amplification technique, which was inspired by work of Bun et al. [BUV14] in the context of differential privacy, in Section 3.2.2). This technique says that, given  $v = k \cdot \log(d/k)$  databases  $\mathcal{D}'_1, \dots, \mathcal{D}'_v$ , each with  $d + k' - 1$  columns and  $1/\epsilon^2$  rows, we can construct a bigger database  $\mathcal{D}$  such that one can use  $\epsilon$ -approximate answers to all  $(k + k')$ -way marginal queries on  $\mathcal{D}$  to obtain  $\epsilon$ -approximate answers to the  $k'$ -way marginals on every database  $\mathcal{D}'_i$ .

This amplification technique actually requires the  $\tilde{\Omega}(d/\epsilon^2)$  lower bound for  $k'$ -way marginals to hold even if the answer vector  $y$  only has error  $\epsilon$  “on average”, rather than having error at most  $\epsilon$  for every single answer. Hence, we cannot directly use the KRSU lower bound in our argument. In fact, to reconstruct a database from answers that have error at most  $\epsilon$  only “on average”, one cannot use  $L_2$  distance minimization as in KRSU’s lower bound argument, since  $L_2$ -minimization is highly sensitive to a few answers having large error. Fortunately, De [De12] shows how to use  $L_1$ -minimization to establish an  $\tilde{\Omega}(1/\epsilon^2)$  lower bound even in the setting in which answers are only required to have error at most  $\epsilon$  “on average”. We use his techniques to obtain a lower bound suitable for our argument.

#### 4.1.2 Proof Details

In the context of differential privacy, De [De12], building on [KRSU10, Rud12], described an algorithm for reconstructing a database  $\mathcal{D}$ , given sufficiently accurate answers to all  $k$ -itemset frequency queries on  $\mathcal{D}$ . In our terminology, De’s result establishes that any For-All-Itemset-Frequency-Estimator-sketch can be used to losslessly encode  $\Omega(\frac{d}{\epsilon^2 \log_{(q)}(1/\epsilon)})$  bits of information. Here  $\log_{(q)}(\cdot)$  denotes the logarithm function iterated  $q$  times. Formally, we use the following slight refinement of De’s result.

**Lemma 20** (Variant of Theorem 5.12 of [De12]). *For any constant integers  $k \geq 2$  and  $q \geq 1$ , there exists a constant  $\gamma = \gamma(k, q) > 0$  and a distribution  $\mu$  over  $k$ -itemset queries such that the following holds.*

*Let  $d$  and  $\epsilon$  be parameters satisfying  $1/\epsilon^2 \leq d^{k-1}/\log_{(q)}(1/\epsilon^2)$ . Suppose  $\mathcal{S}$  is any summary algorithm that can answer a  $1 - \gamma$  fraction of all  $k$ -itemset frequency queries under  $\mu$  on databases with  $d$  columns to error  $\pm\epsilon$ . Then there exists a  $b = b(d, \epsilon) = \Omega(d/\epsilon^2 \log_{(q)}(1/\epsilon))$ , an  $n = n(d, \epsilon) = O(\log_{(q)}(d)/\epsilon^2)$ , a database-generation algorithm  $\mathcal{A}$  that takes as input a Boolean vector  $y \in \{0, 1\}^b$  and outputs a database  $\mathcal{A}(y) \in (\{0, 1\}^d)^n$ , and a decoding algorithm  $\mathcal{A}'$  such that  $\mathcal{A}'$  outputs  $y$  with high probability given  $\mathcal{S}(\mathcal{A}(y))$ .*

We prove Lemma 20 in Appendix B.

Lemma 20 alone is enough to yield a lower bound of  $\Omega\left(d/\epsilon^2 \log_{(q)}(1/\epsilon)\right)$  on the size of For-All-Itemset-Frequency-Estimator sketches capable of answering all  $k$ -itemset frequency queries to error  $\pm\epsilon$ , for any  $k \geq 2$ . The technical contribution of this section is to “bootstrap” this result to obtain a lower bound of  $\Omega\left(kd \log(d/k)/\epsilon^2 \log_{(q)}(1/\epsilon)\right)$  bits, which improves over the bound that follows from a direct application of Lemma 20 even for  $k = 3$ . This lower bound is essentially optimal, matching the  $O(kd \log(d/k)/\epsilon^2)$  upper bound achieved by algorithm SUBSAMPLE up to a  $\log_{(q)}(1/\epsilon)$  factor (for an arbitrarily large constant  $q$ ).

*Proof of Theorem 16.* Let  $v = (k - c) \log(d/(k - c))$  and  $x_1, \dots, x_v \in \{0, 1\}^d$  be the strings promised by Fact 18 applied with  $k' = k - c$ . Recall that for every vector  $s \in \{0, 1\}^v$ , there is a  $(k - c)$ -itemset  $T_s$  such that  $f_{T_s}(x_i) = s_i$  for all  $i \in [v]$ . Recall that  $c \geq 2$  is a parameter of the theorem.

Let  $\gamma = \gamma(c, q)$  be the constant in Lemma 20. Fix  $\epsilon' = 100\epsilon/\gamma$ . Suppose that we are given  $v$  strings  $y_1, \dots, y_v \in \{0, 1\}^b$ , where  $b = b(d, \epsilon')$  is as in Lemma 20 for  $k = c$ . Let  $\mathcal{A}(y_i) = \mathcal{D}_i \in (\{0, 1\}^d)^n$ ,

where  $\mathcal{A}$  is the database generation algorithm promised by Lemma 20. We show how to construct a single “large” database  $\mathcal{D}$  with  $2d$  columns and  $nv$  rows such that 96% of the bits of  $y_1, \dots, y_v$  can be recovered from  $\mathcal{S}(\mathcal{D}, k, \epsilon, \delta)$ . Note that Lemma 20 applies, since  $c \geq 2$ . Also note that  $\gamma$  is indeed a constant since we required that  $c$  and  $q$  are constants; therefore  $\epsilon' = O(\epsilon)$ .

**Definition of  $\mathcal{D}$ .** Recall that  $\mathcal{D}_i(j)$  denotes the  $j$ th row of  $\mathcal{D}_i$ . Define  $\mathcal{D}'_i$  to be the database with  $2d$  columns and  $n$  rows defined via  $\mathcal{D}'_i(j) = (x_i, \mathcal{D}_i(j))$ . That is,  $\mathcal{D}'_i$  is obtained from  $\mathcal{D}_i$  by appending the string  $x_i$  to the front of every row. We define  $\mathcal{D}$  to be the concatenation of all of the  $\mathcal{D}'_i$  databases. We index the  $vn$  rows of  $\mathcal{D}$  as  $(i, j) : i \in [v], j \in [n]$ .

**Reconstructing  $y_1, \dots, y_v$  from  $\mathcal{S}(\mathcal{D})$ .** For any  $c$ -itemset query  $T \subseteq [d]$ , let  $z_T$  denote the vector  $z_T = (f_T(\mathcal{D}_1), \dots, f_T(\mathcal{D}_v))$ . Let  $s \in \{0, 1\}^v$  be any vector. Define  $T' = T'(T, s) \subseteq [2d]$  to be the  $k$ -itemset whose indicator vector is the concatenation of the indicator vectors of  $T_s$  and  $T$ ; that is,  $T' := T_s \cup \{j + d : j \in T\}$ .

We claim that  $\frac{1}{v} \langle s, z_T \rangle = f_{T'(s, T)}(\mathcal{D})$ . To see this, note that  $T'$  is contained in row  $(i, j)$  of  $\mathcal{D}$  if and only if  $s_i = 1$  and  $T$  is contained in row  $j$  of  $\mathcal{D}_i$ . Hence,

$$\frac{1}{v} \langle s, z_T \rangle = \frac{1}{v} \cdot \sum_{i \in [v] : s_i = 1} f_T(\mathcal{D}_i) \quad (6)$$

$$= \frac{1}{v} \sum_{i \in [v]} \frac{1}{n} |\{j \in [n] : T \text{ is contained in } \mathcal{D}_i(j)\}| \quad (7)$$

$$= \frac{1}{nv} \sum_{(i, j) \in [v] \times [n]} |\{T' \text{ is contained in row } (i, j) \text{ of } \mathcal{D}\}| \quad (8)$$

$$= f_{T'(T, s)}(\mathcal{D}). \quad (9)$$

Hence, from any For-All-Itemset-Frequency-Estimator sketch  $\mathcal{S}(\mathcal{D}, k, \epsilon, \delta)$ , one can compute for every  $c$ -itemset  $T$ , an estimate  $\hat{f}_{T'(T, s)}$  satisfying  $|\hat{f}_{T'(T, s)} - \frac{1}{v} \langle s, z_T \rangle| \leq \epsilon$ . The following lemma describes why these estimates are useful in reconstructing  $y_1, \dots, y_v$ .

**Lemma 21.** Fix a  $c$ -itemset  $T \subset [d]$ . Given values  $\hat{f}_{T'(s, T)}$  satisfying  $|\hat{f}_{T'(s, T)} - \frac{1}{v} \langle s, z_T \rangle| \leq \epsilon$  for all  $s \in \{0, 1\}^v$ , it is possible to identify a vector  $\hat{z}_T \in [0, 1]^v$  satisfying  $\frac{1}{v} \|\hat{z}_T - z_T\|_1 \leq 4\epsilon$ .

*Proof.* Consider the algorithm that outputs any vector  $\hat{z}_T \in [0, 1]^v$  satisfying the following property:

$$\text{For all } s \in \{0, 1\}^v, \left| \frac{1}{v} \langle \hat{z}_T, s \rangle - \hat{f}_{T'(s, T)} \right| \leq \epsilon. \quad (10)$$

Note that at least one such vector always exists, because setting  $\hat{z}_T = z_T$  satisfies Equation (10). Thus, the algorithm always produces some output.

We claim that any  $\hat{z}_T$  output by the algorithm satisfies  $\frac{1}{v} \|\hat{z}_T - z_T\|_1 \leq 4\epsilon$ . Indeed, suppose otherwise. Define  $s^{(1)} \in \{0, 1\}^v$  via  $s_i^{(1)} = 1$  if and only if  $\hat{z}_{T,i} \geq z_{T,i}$  and  $s^{(2)} \in \{0, 1\}^v$  via  $s_i^{(2)} = 1$  if and only if  $z_{T,i} > \hat{z}_{T,i}$ . Then either  $\frac{1}{v} (\langle \hat{z}_T, s^{(1)} \rangle - \langle z_T, s^{(1)} \rangle) > 2\epsilon$ , or  $\frac{1}{v} (\langle \hat{z}_T, s^{(2)} \rangle - \langle z_T, s^{(2)} \rangle) > 2\epsilon$ . Assume without loss of generality that the former case holds. Then

$$\left| \frac{1}{v} \langle \hat{z}_T, s^{(1)} \rangle - \hat{f}_{T'(s^{(1)}, T)} \right| \geq \frac{1}{v} (\langle \hat{z}_T, s^{(1)} \rangle - \langle z_T, s^{(1)} \rangle) - \left| \frac{1}{v} \langle z_T, s \rangle - \hat{f}_{T'(s, T)} \right| > 2\epsilon - \epsilon = \epsilon,$$

where the inequality holds by the triangle inequality. But this contradicts the assumption that  $\hat{z}_T$  satisfies (10).  $\square$

For each  $c$ -itemset  $T$ , let  $\hat{z}_T$  be as in Lemma 21. We think of  $\hat{z}_{T,i}$  as an estimate of  $z_{T,i} = f_T(\mathcal{D}_i)$ . Lemma 21 guarantees that *for any distribution  $\mu$  over  $c$ -itemsets  $T$* , this estimate has error at most  $4\epsilon$  on average, when the averaging is done over a random  $c$ -itemset  $T$  chosen according to  $\mu$ , and databases  $\mathcal{D}_i$ . In symbols:

$$\mathbf{E}_{T \leftarrow \mu} \mathbf{E}_{i \in v} |\hat{z}_{T,i} - z_{T,i}| = \sum_T \mu(T) (\mathbf{E}_{i \in v} |\hat{z}_{T,i} - z_{T,i}|) \leq \sum_T \mu(T) \cdot 4\epsilon \leq 4\epsilon.$$

Here,  $\mathbf{E}_{T \leftarrow \mu}$  denotes the expectation operation when  $T$  is chosen according to the distribution  $\mu$ , and the penultimate inequality holds by Lemma 21.

By Markov's inequality we conclude that for at least 96% of the databases  $\mathcal{D}_i$ , the estimates  $\hat{z}_{T,i}$  have error at most  $100\epsilon$  on average, where the averaging is over the choice of  $T$  according to distribution  $\mu$ . That is, for at least 96% of databases  $\mathcal{D}_i$ , it holds that

$$\mathbf{E}_{T \leftarrow \mu} \mathbf{E}_{i \in v} |\hat{z}_{T,i} - z_{T,i}| \leq 100\epsilon. \quad (11)$$

For any  $\mathcal{D}_i$  satisfying Equation (11) and any  $\gamma > 0$ , another application of Markov's inequality implies that the  $|\hat{z}_{T,i} - f_T(\mathcal{D}_i)| \leq 100\epsilon/\gamma = \epsilon'$  for a  $1 - \gamma$  fraction of all  $c$ -itemsets  $T$  under distribution  $\mu$ . By Lemma 20, this implies that  $y_i$  can be exactly recovered from the  $\hat{z}_{T,i}$  values, using algorithm  $\mathcal{A}'$ .

Since 96% of the  $y_i$  vectors can be *exactly* recovered, it follows that at least 96% of the  $bv$  total bits in the vectors of  $y_1, \dots, y_v$  can be recovered. Suppose we let the  $bv$  bits in the collection of vectors  $(y_1, \dots, y_v)$  be the error-corrected encoding of a single vector  $(y'_1, \dots, y'_z) \in \{0, 1\}^z$ , using a code with constant rate that is uniquely decodable from 4% errors (e.g. using a Justesen code [Jus72]). Then  $z = \Omega(bv)$ , and it follows from the above that  $(y'_1, \dots, y'_z)$  can be *exactly* reconstructed from  $\mathcal{S}(\mathcal{D}, k, \epsilon, \delta)$  with probability  $1 - \delta$ . Basic information theory then implies that

$$|\mathcal{S}(\mathcal{D}, k, \epsilon, \delta)| = \Omega(bv) = \Omega\left(\frac{kd \log(d/k)}{(\epsilon')^2 \log_{(q)}(1/(\epsilon'))}\right) = \Omega\left(\frac{kd \log(d/k)}{\epsilon^2 \log_{(q)}(1/\epsilon)}\right),$$

where we have used the fact that  $\epsilon' = O(\epsilon)$ . This completes the proof of the theorem.  $\square$

## 4.2 The For-Each Case: Proof of Theorem 17

Recall that Theorem 17 establishes a lower bound against For-Each-Itemset-Frequency-Estimator sketches that is tight up to a  $\log_{(q)}(d)$  factor. We prove Theorem 17 via a simple argument that shows how to transform any For-Each-Itemset-Frequency-Estimator sketch into a For-All-Itemset-Frequency-Estimator sketch with a modest increase in space. This allows us to transform Theorem 16 into the claimed lower bound against For-Each-Itemset-Frequency-Estimator sketches.

*Proof of Theorem 17.* Suppose that we are given an For-Each-Itemset-Frequency-Estimator sketching algorithm  $\mathcal{S}$  using space  $|\mathcal{S}|$ . We show how to transform  $\mathcal{S}$  into a For-All-Itemset-Frequency-Estimator sketching algorithm  $\mathcal{S}'$  using space  $O\left(|\mathcal{S}| \cdot \log\left(\frac{d}{k}\right)\right) = O(|\mathcal{S}(k, \epsilon, \delta)| \cdot k \cdot \log(d/k))$ . It then follows from Theorem 16 that  $|\mathcal{S}| = \Omega\left(\frac{d}{\epsilon^2 \log_{(q)}(1/\epsilon)}\right)$ .

The For-All-Itemset-Frequency-Estimator sketching algorithm  $\mathcal{S}'$  simply outputs  $10 \cdot \log\left(\frac{d}{k}/\delta\right)$  independent copies of  $\mathcal{S}(\mathcal{D})$  (i.e., using fresh randomness for each of the  $10 \cdot \log\left(\frac{d}{k}/\delta\right)$  runs of  $\mathcal{S}$ ). Given any  $k$ -itemset  $T$ , the query procedure  $\mathcal{Q}'$  associated with  $\mathcal{S}'$  simply runs the query procedure  $\mathcal{Q}$  associated with  $\mathcal{S}$  on each of the copies of  $\mathcal{S}(\mathcal{D})$ , and outputs the median of the results. Since each copy of  $\mathcal{S}$  outputs an estimated frequency that has error at most  $\epsilon$  with probability  $1 - \delta > 1/2$ , standard Chernoff Bounds

imply that for any fixed  $k$ -itemset  $T$ , the median estimate will have error at most  $\epsilon$  with probability at least  $1 - \delta / \binom{d}{k}$ . A union bound implies that the median estimate will have error at most  $\epsilon$  for all  $\binom{d}{k}$  itemsets with probability at least  $1 - \delta$ . Thus,  $\mathcal{S}'$  is a For-All-Itemset-Frequency-Estimator sketching algorithm with failure probability at most  $\delta$ .  $\square$

## 5 Conclusion

We introduced four closely related notions that capture the problem of approximating itemset frequencies in databases. For all four problems, we studied the minimal size of sketches that permit a user to recover sufficiently accurate information about itemset frequencies. After identifying three naive algorithms that apply to all four problems, we turned to proving sketch size lower bounds. Our results establish that random sampling achieves optimal or essentially optimal sketch size for all four problems. This stands in contrast to several seemingly similar problems, such as identifying approximate frequent items in data streams, and various matrix approximation problems, for which uniform sampling is not an optimal sketching algorithm.

We proved our lower bounds by adapting and extending techniques developed in the literature on differentially privacy data analysis. It is an interesting open question whether there are other problems in non-private data analysis that can be resolved using techniques from the literature on differential privacy.

In addition, our lower bound arguments specify a “hard” distribution over databases, for which it is impossible to improve upon the space usage of the uniform sampling sketching algorithm for answering approximate itemset frequency queries. But real-world databases are likely to be substantially more structured than the databases appearing in our hard distribution, and real-world query loads are likely to be highly non-uniform. In these settings, *importance sampling* is a natural candidate for improving upon the space usage of the uniform sampling sketching algorithm. It would be interesting to identify rigorous yet realistic conditions on databases and query loads that allow for such an improvement. Subsequent work by Lang et al. [LLS16] takes some initial steps in this direction.

**Acknowledgements.** The authors are grateful to Amit Chakrabarti, Graham Cormode, Nikhil Srivastava, and Suresh Venkatasubramanian for several helpful conversations during the early stages of this work.

## References

- [Abl96] Farid M. Ablayev. Lower bounds for one-way probabilistic communication complexity and their application to space complexity. *Theor. Comput. Sci.*, 157(2):139–159, 1996.
- [AIS93] Rakesh Agrawal, Tomasz Imieliński, and Arun Swami. Mining association rules between sets of items in large databases. *SIGMOD Rec.*, 22(2):207–216, 1993.
- [AKW14] Alexandr Andoni, Robert Krauthgamer, and David P. Woodruff. The sketching complexity of graph cuts. *CoRR*, abs/1403.7058, 2014.
- [BBR03] Jean-François Boulcaut, Artur Bykowski, and Christophe Rigotti. Free-sets: A condensed representation of boolean data for the approximation of frequency queries. *Data Min. Knowl. Discov.*, 7(1):5–22, 2003.
- [BCD<sup>+</sup>07] Boaz Barak, Kamalika Chaudhuri, Cynthia Dwork, Satyen Kale, Frank McSherry, and Kunal Talwar. Privacy, accuracy, and consistency too: a holistic solution to contingency table release. In Leonid Libkin, editor, *PODS*, pages 273–282. ACM, 2007.

- [BUV14] Mark Bun, Jonathan Ullman, and Salil Vadhan. Fingerprinting codes and the price of approximate differential privacy. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing, STOC '14*, pages 1–10, New York, NY, USA, 2014. ACM.
- [CG07] Toon Calders and Bart Goethals. Non-derivable itemset mining. *Data Min. Knowl. Discov.*, 14(1):171–206, February 2007.
- [CKN08] James Cheng, Yiping Ke, and Wilfred Ng. A survey on algorithms for mining frequent itemsets over data streams. *Knowl. Inf. Syst.*, 16(1):1–27, 2008.
- [CTUW14] Karthekeyan Chandrasekaran, Justin Thaler, Jonathan Ullman, and Andrew Wan. Faster private release of marginals on small databases. In Moni Naor, editor, *ITCS*, pages 387–402. ACM, 2014.
- [De12] Anindya De. Lower bounds in differential privacy. In Ronald Cramer, editor, *TCC*, volume 7194 of *Lecture Notes in Computer Science*, pages 321–338. Springer, 2012.
- [DMNS06] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography, Third Theory of Cryptography Conference, TCC 2006, New York, NY, USA, March 4-7, 2006, Proceedings*, pages 265–284, 2006.
- [FK04] Uriel Feige and Shimon Kogan. Hardness of approximation of the Balanced Complete Bipartite Subgraph problem. Technical Report MCS04-04, Weizmann Institute of Science, 2004.
- [GHRU13] Anupam Gupta, Moritz Hardt, Aaron Roth, and Jonathan Ullman. Privately releasing conjunctions and the statistical query barrier. *SIAM J. Comput.*, 42(4):1494–1520, 2013.
- [GWWZ15] Dirk Van Gucht, Ryan Williams, David P. Woodruff, and Qin Zhang. The communication complexity of distributed set-joins with applications to matrix multiplication. In Tova Milo and Diego Calvanese, editors, *Proceedings of the 34th ACM Symposium on Principles of Database Systems, PODS 2015, Melbourne, Victoria, Australia, May 31 - June 4, 2015*, pages 199–212. ACM, 2015.
- [HCW06] Matthew Hamilton, Rhonda Chaytor, and Todd Wareham. The parameterized complexity of enumerating frequent itemsets. In *Proceedings of the Second International Conference on Parameterized and Exact Computation, IWPEC'06*, pages 227–238, Berlin, Heidelberg, 2006. Springer-Verlag.
- [Jus72] J. Justesen. A class of constructive asymptotically good algebraic codes. *IEEE Transactions on Information Theory*, 18(5):652–656, 1972.
- [KRSU10] Shiva Prasad Kasiviswanathan, Mark Rudelson, Adam Smith, and Jonathan Ullman. The price of privately releasing contingency tables and the spectra of random matrices with correlated rows. In *STOC*, pages 775–784, 2010.
- [LLS16] Kevin Lang, Edo Liberty, and Konstantin Shmakov. Stratified sampling meets machine learning, 2016. Manuscript. Available at <http://www.cs.yale.edu/homes/el327/papers/lls15.pdf>.
- [LLSW05] Jinyan Li, Haiquan Li, Donny Soh, and Limsoon Wong. A correspondence between maximal complete bipartite subgraphs and closed patterns. In *PKDD, PKDD'05*, pages 146–156, Berlin, Heidelberg, 2005. Springer-Verlag.



- [LMT14] Edo Liberty, Michael Mitzenmacher, and Justin Thaler. Space lower bounds for itemset frequency sketches. *CoRR*, abs/1407.3740, 2014.
- [MM02] Gurmeet Singh Manku and Rajeev Motwani. Approximate frequency counts over data streams. In *Proceedings of the 28th International Conference on Very Large Data Bases, VLDB '02*, pages 346–357. VLDB Endowment, 2002.
- [MT96] Heikki Mannila and Hannu Toivonen. Multiple uses of frequent sets and condensed representations. In *KDD*, 1996.
- [MT07] Frank McSherry and Kunal Talwar. Mechanism design via differential privacy. In *Proceedings of the 48th Annual IEEE Symposium on Foundations of Computer Science, FOCS '07*, pages 94–103, Washington, DC, USA, 2007. IEEE Computer Society.
- [PDZH04] Jian Pei, Guozhu Dong, Wei Zou, and Jiawei Han. Mining condensed frequent-pattern bases. *Knowl. Inf. Syst.*, 6(5):570–594, 2004.
- [Pri14] Eric Price. Optimal lower bound for itemset frequency indicator sketches. *CoRR*, abs/1410.2640, 2014.
- [Rud12] Mark Rudelson. Row products of random matrices. *Advances in Mathematics*, 231(6):3199 – 3231, 2012.
- [TUV12] Justin Thaler, Jonathan Ullman, and Salil P. Vadhan. Faster algorithms for privately releasing marginals. In Artur Czumaj, Kurt Mehlhorn, Andrew M. Pitts, and Roger Wattenhofer, editors, *ICALP (1)*, volume 7391 of *Lecture Notes in Computer Science*, pages 810–821. Springer, 2012.
- [Yan04] Guizhen Yang. The complexity of mining maximal frequent itemsets and maximal frequent patterns. In *KDD*, KDD '04, pages 344–353, New York, NY, USA, 2004. ACM.

## A Proof of Fact 18

We provide a direct construction of the vectors whose existence is guaranteed by Fact 18, restated here for convenience.

**Fact 18.** *For any  $k' \geq 1$ , let  $v = k' \cdot \log(d/k')$ . There exist strings  $x_1, \dots, x_v \in \{0, 1\}^d$  such that for every string  $s \in \{0, 1\}^v$ , there is a  $k'$ -itemset  $T_s$  such that  $f_{T_s}(x_i) = s_i$  for all  $i \in [v]$ .*

*Proof of Fact 18.* For expository purposes, we first describe a set of vectors  $w_1, \dots, w_{k'} \in \{0, 1\}^{k'}$  that are “shattered” by  $k'$ -itemset frequency queries, i.e., for every string  $s \in \{0, 1\}^{k'}$ , there is a  $k'$ -itemset  $T_s$  such that  $f_{T_s}(y_i) = s_i$  for all  $i \in [k']$ . We then describe a set of vectors  $y_1, \dots, y_{\log d} \in \{0, 1\}^d$  that are shattered even by 1-itemset frequency queries. Finally, we explain how to “glue together” the  $w_i$ ’s and  $y_i$ ’s to obtain the full set  $x_1, \dots, x_v \in \{0, 1\}^d$  whose existence is claimed in the statement of Fact 18.

**Description of the  $w_i$ ’s.** For each  $i \in [k']$ , define  $w_i \in \{0, 1\}^{k'}$  via:

$$\begin{cases} w_{i,j} = 1 & : \quad 1 \leq j \leq k', j \neq i \\ w_{i,j} = 0 & \quad j = i \end{cases}$$

To restate the above in matrix notation, we define the  $k' \times k'$  matrix  $W^{(k')}$  whose rows are the  $w_i$ ’s via:

$$W^{(k')} := \begin{pmatrix} w_1 \\ w_2 \\ \vdots \\ w_{k'-1} \\ w_{k'} \end{pmatrix} = \begin{pmatrix} 0 & 1 & 1 & \dots & 1 & 1 \\ 1 & 0 & 1 & \dots & 1 & 1 \\ \vdots & \vdots & \vdots & \dots & \vdots & \vdots \\ 1 & 1 & 1 & \dots & 0 & 1 \\ 1 & 1 & 1 & \dots & 1 & 0 \end{pmatrix}$$

For any string  $s \in \{0, 1\}^{k'}$ , let  $T_s := \{i : s_i = 0\}$ . It is straightforward to check that  $f_{T_s}(w_i) = s_i$  as desired.

**Description of the  $y_i$ 's.** For each  $i \in [\log d]$ , we define each  $y_i \in \{0, 1\}^d$  to ensure that the matrix whose rows are the  $y_i$ 's contains every possible  $\log(d)$ -bit string as a column. In matrix notation, we define the  $\log(d) \times d$  matrix  $Y^{(d)}$  via:

$$Y^{(d)} := \begin{pmatrix} y_1 \\ y_2 \\ \vdots \\ y_{\log(d)-1} \\ y_{\log d} \end{pmatrix} = \begin{pmatrix} 0 & 0 & 0 & \dots & 1 & 1 \\ 0 & 0 & 0 & \dots & 1 & 1 \\ \vdots & \vdots & \vdots & \dots & \vdots & \vdots \\ 0 & 0 & 1 & \dots & 1 & 1 \\ 0 & 1 & 0 & \dots & 0 & 1 \end{pmatrix}$$

For any string  $s \in \{0, 1\}^{\log d}$ , we interpret  $s$  as the binary representation of an integer  $\text{int}(s) \in \{0, \dots, d-1\}$ , and define  $T_s := \{\text{int}(s)\}$ . It is straightforward to check that  $f_{T_s}(y_i) = s_i$  as desired.

**Description of the  $x_i$ 's.** Recall that  $v = k' \cdot \log(d/k')$ . Consider the  $v \times d$  matrix  $X$  whose rows are the  $x_i$ 's. We view this matrix as a collection of sub-matrices, where each sub-matrix has dimension  $k' \times (d/k')$ . More specifically, let  $\mathbf{J}$  denote the  $k' \times (d/k')$  matrix of all-ones. We define  $X$  to be the matrix obtained from  $W^{(d/k')}$  by replacing each entry of  $W^{(d/k')}$  that is equal to 1 with the matrix  $\mathbf{J}$ , and replacing each entry of  $W^{(d/k')}$  that is equal to 0 with the matrix  $Y^{(d/k')}$ . In more detail, define:

$$X := \begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_{k'-1} \\ x_{k'} \end{pmatrix} = \begin{pmatrix} Y^{(d/k')} & \mathbf{J} & \mathbf{J} & \dots & \mathbf{J} & \mathbf{J} \\ \mathbf{J} & Y^{(d/k')} & \mathbf{J} & \dots & \mathbf{J} & \mathbf{J} \\ \vdots & \vdots & \vdots & \dots & \vdots & \vdots \\ \mathbf{J} & \mathbf{J} & \mathbf{J} & \dots & Y^{(d/k')} & \mathbf{J} \\ \mathbf{J} & \mathbf{J} & \mathbf{J} & \dots & \mathbf{J} & Y^{(d/k')} \end{pmatrix}$$

Given any vector  $s \in \{0, 1\}^v$ , we interpret  $s$  as specifying  $k'$  integers  $\ell_1, \dots, \ell_{k'} \in \{0, \dots, d/k'\}$  in the natural way. We view  $[d]$  as the cross-product  $[k'] \times [d/k']$ , and associate each  $j \in [d]$  with a pair  $(r_1, r_2) \in [k'] \times [d/k']$  in the natural way. We then define  $T_s := \{(i, \ell_i) : i \in [k']\}$ . It is then straightforward to observe that  $f_{T_s}(x_i) = s_i$ .  $\square$

## B Proof of Lemma 20

We restate Lemma 20 for convenience, before providing its proof.

**Lemma 20.** (Refinement of Theorem 5.12 of [De12]) *For any constant integers  $k \geq 2$  and  $q \geq 1$ , there exists a constant  $\gamma = \gamma(k, q) > 0$  and a distribution  $\mu$  over  $k$ -itemset queries such that the following holds.*

*Let  $d$  and  $\epsilon$  be parameters satisfying  $1/\epsilon^2 \leq d^{k-1}/\log_{(q)}(1/\epsilon^2)$ . Suppose  $\mathcal{S}$  is any summary algorithm that can answer a  $1 - \gamma$  fraction of all  $k$ -itemset frequency queries under  $\mu$  on databases with  $d$  columns to error  $\pm\epsilon$ . Then there exists a  $b = b(d, \epsilon) = \Omega(d/\epsilon^2 \log_{(q)}(1/\epsilon))$ , an  $n = n(d, \epsilon) = O(\log_{(q)}(d)/\epsilon^2)$ , a*

database-generation algorithm  $\mathcal{A}$  that takes as input a Boolean vector  $y \in \{0, 1\}^b$  and outputs a database  $\mathcal{A}(y) \in (\{0, 1\}^d)^n$ , and a decoding algorithm  $\mathcal{A}'$  such that  $\mathcal{A}'$  outputs  $y$  with high probability given  $\mathcal{S}(\mathcal{A}(y))$ .

As Lemma 20 is a refinement of Theorem 5.12 of De's work [De12], the presentation of our proof borrows heavily from De's.

*Proof.* We begin by defining the Hadamard product of matrices.

**Definition 22** (Hadamard product of matrices). *Let  $A_1, \dots, A_s \in \mathbb{R}^{\ell_i \times n}$ . Then, the Hadamard product of  $A_1, \dots, A_s$  is denoted by  $A = A_1 \circ A_2 \circ \dots \circ A_s \in \mathbb{R}^{L \times n}$ , where  $L = \ell_1 \cdot \dots \cdot \ell_s$  and is defined as follows: Every row of  $A$  is identified with a unique element of  $[\ell_1] \times \dots \times [\ell_s]$ . For  $i = (i_1, \dots, i_s)$ , define*

$$A[i, h] = \prod_{j=1}^s A_j[i_j, h]$$

where  $A[i, h]$  represents the element in row  $i$  and column  $h$  of  $A$ .

We will also require the definition of Euclidean sections, which play an important role in the analysis of LP decoding algorithms.

**Definition 23** (Euclidean Sections).  *$V \subseteq \mathbb{R}^z$  is said to be a  $(\delta, d', z)$  Euclidean Section if  $V$  is a linear subspace of dimension  $d'$  and for every  $x \in V$ , the following holds:*

$$\sqrt{z}\|x\|_2 \geq \|x\|_1 \geq \delta\sqrt{z}\|x\|_2.$$

A linear operator  $A: \mathbb{R}^{d'} \rightarrow \mathbb{R}^z$  is said to be  $\delta$ -Euclidean if the range of  $A$  is a Euclidean  $(\delta, d', z)$  section.

The following lemma follows directly from the proof of [De12, Lemma 5.9].

**Lemma 24** (Reformulation of Lemma 5.9 of [De12]). *Let  $d_0^{k-1} > n$ . Suppose there exist Boolean matrices  $A_1, \dots, A_{\ell-1} \in \mathbb{R}^{d_0 \times n}$  such that  $A = A_1 \circ A_2 \circ \dots \circ A_{\ell-1}$ , all the singular values of  $A$  are at least  $\sigma$ , and the range of  $A$  is a  $(\delta, n, d_0^{k-1})$ -Euclidean section.*

*Let  $\mathcal{D}_0$  denote the database with  $n$  rows and  $(k-1) \cdot d_0$  columns obtained from the  $A_i$ 's as follows: the  $j$ th row of  $\mathcal{D}_0$  is the concatenation of the  $j$ th row of each of the matrices  $A_1^T, A_2^T, \dots, A_{\ell-1}^T$ . Let  $\mathcal{A}_1$  denote the database generation algorithm that takes as input a Boolean vector  $y \in \{0, 1\}^n$ , and outputs the database  $\mathcal{D}_1(y)$  with  $n$  rows and  $d_1 := (k-1) \cdot d_0 + 1$  columns obtained from  $\mathcal{D}_0$  by appending an additional column equal to  $y$ .*

*Then, there exists a constant  $\gamma_1 = \gamma_1(\delta) > 0$ , a distribution  $\mu_1$  over  $k$ -itemsets  $T \subseteq [d_1]$ , and a reconstruction algorithm  $\mathcal{R}$  satisfying the following. Fix any  $\zeta_1 \in o(\sqrt{n}\sigma/\sqrt{d_1^{k-1}})$ . Suppose  $\mathcal{R}$  is given  $\mathcal{D}_0$  and approximate itemset frequencies  $\hat{f}_T$  for all  $k$ -itemsets  $T \subseteq [d_1]$ . Let  $S_1$  denote the set of all  $k$ -itemsets  $T$  satisfying  $n \cdot |\hat{f}_T - f_T(\mathcal{D}_2(y))| \leq \zeta_1$ , and suppose that  $\sum_{T \in S_1} \mu(T) \geq 1 - \gamma_1$ . Then  $\mathcal{R}$  outputs a vector  $\hat{y}$  of Hamming distance  $o(n)$  from  $y$ .*

We use Lemma 24 to establish the following stronger statement.

**Lemma 25.** *Let  $A_1, \dots, A_{\ell-1}, \delta, \mathcal{D}_0, \mu_1, d_0, d_1, n, \gamma_1$ , and  $\zeta_1$  be as in Lemma 24. There is a  $b = b(d_0, n) \in \Omega(d_0 \cdot n)$  and database generation algorithm  $\mathcal{A}_2$  that takes as input a Boolean vector  $y' \in \{0, 1\}^b$ , and outputs a database  $\mathcal{D}_2$  with  $n$  rows and  $d_2 := (k-1) \cdot d_0 + d_0 = k \cdot d_0$  columns such that the following holds.*

*There exists a constant  $\gamma_2 = \gamma_2(\delta) > 0$ , a distribution  $\mu_2$  over  $k$ -itemsets  $T \subseteq [d_2]$ , and a reconstruction algorithm  $\mathcal{R}_2$  satisfying the following. Suppose  $\mathcal{R}_2$  is given  $\mathcal{D}_0$  and approximate itemset frequencies  $\hat{f}_T$  for all  $k$ -itemsets  $T \subseteq [d_2]$ . Let  $S_2$  denote the set of all  $k$ -itemsets  $T$  satisfying  $n \cdot |\hat{f}_T - f_T(\mathcal{D}_2(y'))| \leq \zeta_1$ , and suppose that  $\sum_{T \in S_2} \mu_2(T) \geq 1 - \gamma_2$ . Then  $\mathcal{R}_2$  outputs  $y'$ .*

*Proof.* Let  $\mathcal{A}_2$  be the database generation algorithm that takes as input a Boolean vector  $y' \in \{0, 1\}^b$ , and first replaces  $y'$  with an error-corrected encoding  $\text{Enc}(y') \in \{0, 1\}^{d_0 \cdot n}$  of  $y'$ , using an error-correcting code of constant rate that is uniquely decodable from 2% errors.  $\mathcal{A}_2$  then outputs the database  $\mathcal{D}_2(y')$  with  $n$  rows and  $d_2$  columns obtained from  $\mathcal{D}_0$  by appending  $d_0$  additional columns, with the first additional column equal to the first  $n$  bits of  $\text{Enc}(y')$ , the second additional column equal to the second  $n$  bits of  $\text{Enc}(y')$ , and so on. We refer to the  $d_0$  attributes corresponding to these additional columns as *special attributes*. Similarly, we call an itemset  $T \subseteq [d_2]$  *special* if  $T$  contains exactly one special attribute.

For each  $i \in [d_0]$ , let  $y^{(i)} \in \{0, 1\}^n$  denote the vector  $(\text{Enc}(y')_{(i-1) \cdot n+1}, \dots, \text{Enc}(y')_{i \cdot n})$ ; that is,  $y^{(i)}$  is the  $i$ th “block” of  $n$  bits from  $\text{Enc}(y')$ . Let  $\mathcal{D}_1(y^{(i)})$  be as in the statement of Lemma 24. Note that  $\mathcal{D}_1(y^{(i)})$  is a sub-database of  $\mathcal{D}_2(y')$ , in the sense that  $\mathcal{D}_1(y^{(i)})$  equals  $\mathcal{D}_2(y')$  with several columns removed. Hence, for any  $k$ -itemset  $T_1 \subseteq [d_1]$ , there is a unique itemset  $g_i(T_1) \subseteq [d_2]$  such that  $f_{T_1}(\mathcal{D}_1(y^{(i)})) = f_{g_i(T_1)}(\mathcal{D}_2(y'))$ . Notice that  $g_i(T_1)$  is a special itemset, for any  $i$  and  $T_1$ . Moreover, the  $g_i$ ’s are all invertible: for any special itemset  $T_2 \subseteq [d_2]$ , there is a unique itemset  $h(T_2) \subseteq [d_1]$  and a unique  $i$  satisfying  $g_i(h(T_2)) = T_2$ .

Let  $\gamma_2 = \gamma/100$ . We define the distribution  $\mu_2$  over  $k$ -itemsets  $T \subseteq [d_2]$  as follows.

$$\mu_2(T) = \begin{cases} 0 & \text{if } T \text{ is not special.} \\ (1/d_0) \cdot \mu_1(h(T)) & \text{if } T \text{ is special.} \end{cases}$$

As per the hypothesis of the lemma, suppose  $\mathcal{R}_2$  is given  $\mathcal{D}_0$  and approximate itemset frequencies  $\hat{f}_T$  for all  $k$ -itemsets  $T \subseteq [d_2]$ . Let  $S_2$  denote the set of all  $k$ -itemsets  $T$  satisfying  $n \cdot |\hat{f}_T - f_T(\mathcal{D}_2(y'))| \leq \zeta_1$ , and suppose that  $\sum_{T \in S_2} \mu_2(T) \geq 1 - \gamma_2$ .

The recovery algorithm  $\mathcal{R}_2$  will reconstruct  $y'$  by first constructing a vector  $y'' \in \{0, 1\}^{d_0 \cdot n}$  such that the fractional Hamming distance between  $y''$  and  $\text{Enc}(y')$  is at most .02, and then running the decoding algorithm for the error-correcting code on  $y''$ .  $\mathcal{R}_2$  constructs the vector  $y''$  as follows. For each  $i \in [d_0]$ ,  $\mathcal{R}_2$  constructs the  $i$ ’th block of  $n$  bits of  $y''$  by simulating  $\mathcal{R}_1$  on  $\mathcal{D}_1(y^{(i)})$  in the natural way: whenever  $\mathcal{R}_1$  requests a value  $\hat{f}_{T_1}$ ,  $\mathcal{R}_2$  returns the value  $f_{g_i(T_1)}$ .  $\mathcal{R}_2$  then sets  $(y''_{(i-1) \cdot n+1}, \dots, y''_{i \cdot n})$  to the vector  $\hat{y}$  output by  $\mathcal{R}_1$ .

**Showing  $y''$  is close to  $\text{Enc}(y')$  in Hamming distance.** For each special attribute  $i$ , let  $S_{1,i}$  denote the set of itemsets  $T \subseteq [d_2]$  in  $\text{Range}(g_i)$  satisfying  $n \cdot |\hat{f}_T - f_T(\mathcal{D}_2(y'))| \leq \zeta_1$ . Since a  $1 - \gamma_2$  fraction of the estimates  $\hat{f}_T$  under  $\mu_2$  satisfy  $n \cdot |\hat{f}_T - f_T(\mathcal{D}_2(y'))| \leq \zeta_1$ , Markov’s inequality implies that 99% of the  $i$ ’s satisfy  $\sum_{T \in S_{1,i}} \mu_2(T)/d_0 \geq 1 - 100\gamma_2 = 1 - \gamma_1$ . Lemma 24 implies that for each such  $i$ , the  $i$ th block output by  $\mathcal{R}_2$ , namely  $(y''_{(i-1) \cdot n+1}, \dots, y''_{i \cdot n})$ , will have Hamming distance  $o(n)$  from the  $i$ th block of  $\text{Enc}(y')$ . Hence,  $y''$  has fractional Hamming distance at most  $.01 + o(1) \leq .02$  from  $\text{Enc}(y')$ .  $\square$

Rudelson [Rud12] proved the existence of matrices  $A_1, \dots, A_{k-1}$  satisfying the conditions of Lemmas 24 and 25.

**Lemma 26** (Rudelson [Rud12], see also Theorem 5.11 of [De12]). *Let  $q, k$  be constants. Also, let  $\nu \sim \mathbb{R}^{d' \times n}$  be a distribution over matrices such that every entry of the matrix is an independent and unbiased  $\{0, 1\}$  random variable. Let  $A_1, \dots, A_{k-1}$  be i.i.d. copies of random matrices drawn from the distribution  $\nu$  and  $A$  be the Hadamard product of  $A_1, \dots, A_{k-1}$ . Then, provided that  $d^{k-1} = o(n \log_{(q)}(n))$ , with probability  $1 - o(1)$ , the smallest singular value of  $A$ , denoted by  $\sigma_n(A)$ , satisfies  $\sigma_n(A) = \Omega(\sqrt{d^{k-1}})$ . Also, the range of  $A$  is a  $(\gamma(q, \ell), n, d^{k-1})$  Euclidean section for some  $\gamma(q, \ell) > 0$ .*

Combining Lemmas 26 and Lemma 25, we obtain the following lemma.

**Lemma 27.** *For any positive constants  $q, k$ , and any pair  $d_0, n > 0$  satisfying  $n \log_{(q)}(n) < d_0^{k-1}$ , there is a  $b = b(d_0, n) \in \Omega(d_0 \cdot n)$  and database generation algorithm  $\mathcal{A}_2$  that takes as input a Boolean vector  $y \in \{0, 1\}^b$ , and outputs a database  $\mathcal{D}_2$  with  $n$  rows and  $d_2 := (k-1) \cdot d_0 + d_0 = k \cdot d_0$  columns such that the following holds.*

*There exists a constant  $\gamma_2 = \gamma_2(k, q) > 0$ , a distribution  $\mu_2$  over  $k$ -itemsets  $T \subseteq [d_2]$ , and a reconstruction algorithm  $\mathcal{R}_2$  satisfying the following. Let  $\zeta_1 = \sqrt{n} / \log_{(q+1)}(n)$ . Suppose  $\mathcal{R}_2$  is given  $\mathcal{D}_0$  and approximate itemset frequencies  $\hat{f}_T$  for all  $k$ -itemsets  $T \subseteq [d_2]$ . Let  $S_2$  denote the set of all  $k$ -itemsets  $T$  satisfying  $n \cdot |\hat{f}_T - f_T(\mathcal{D}_2(y'))| \leq \zeta_1$ , and suppose that  $\sum_{T \in S_2} \mu_2(T) \geq 1 - \gamma_2$ . Then  $\mathcal{R}_2$  outputs  $y'$ .*

**Remark:** Note that Lemma 27 actually holds for any  $\zeta_1 = o(\sqrt{n})$ ; we choose a particular  $\zeta_1$  that makes the lemma particularly convenient to apply in our context.

For any  $\epsilon > 0$ , suppose we set  $n = 1/(\epsilon^2 \cdot \log_{(q)}(n))$  in the statement of Lemma 27. This causes  $\zeta_1/n$  to equal  $1/(\sqrt{n} \cdot \log_{(q+1)}(n)) = \epsilon \sqrt{\log_{(q)}(n)} / \log_{(q+1)}(n) > \epsilon$ . Hence, we conclude that an For-All-Itemset-Frequency-Estimator sketch that can answer all  $k$ -itemset frequency queries with error bounded by  $\epsilon$  provides sufficiently accurate itemset frequency estimates  $\hat{f}_T$  to apply Lemma 27 with  $n = 1/(\epsilon^2 \cdot \log_{(q)}(n))$ , and Lemma 20 follows. □